

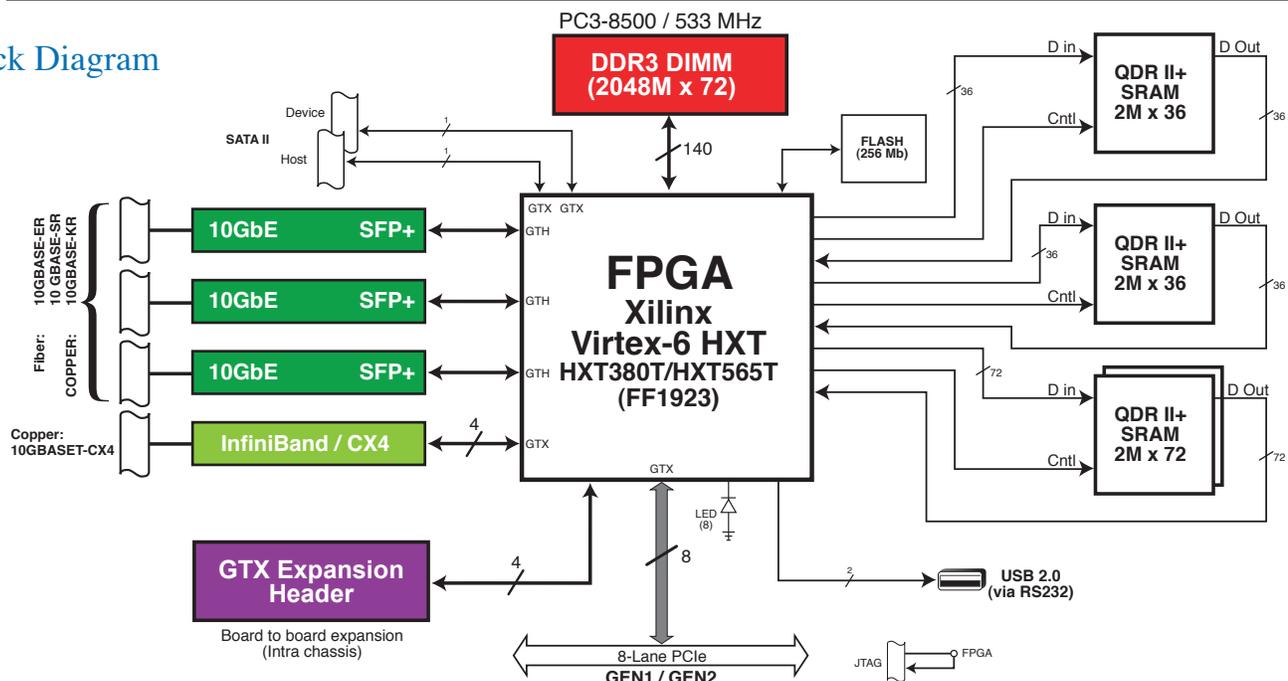
### Features

- 3 separate 10GbE LAN/WAN using SFP+ modules
  - Customized IP for packet analysis with minimum latency
- 1 CX4 socket for low latency InfiniBand networking
- 2 Serial ATA II ports (SATA II) for connection to local hard drive
- Hosted in a an 8-lane GEN1 or GEN2 PCIe slot
  - Stand-alone operations supported with external ATX power supply
- Xilinx Virtex-6 HXT FPGA (FF1923) :
  - HX565T-2,-1 (fastest to slowest)
  - HX380T-3,-2,-1
  - 2M ASIC gates (ASIC measure) when stuffed with Virtex-6 HX565T
    - 354k flip-flop/6-input LUTs (708k total FFs)
    - 4MB total FPGA block memory (1824, 18 kbit blocks)
    - 864, 25x18 multipliers
- Bulk memory: DDR3 DIMM
  - 72-bit data width (64-bit with 8-bit ECC)
  - 533MHz operation, PC3-8500
  - Addressing/power to support 16GB (+ ECC)
- DDR3 Verilog/VHDL reference design provided (no charge)
  - Optimized DDR3 controller for lowest latency bulk memory access
- Optional RLDRAM DIMM instead of DDR3 for ultra low latency
- 3 independent QDR II+ SRAM memory channels:
  - 2 channels: 4M x 36 (144Mb)
    - Separate 36-bit read and write ports
  - 1 channel: 4M x 72 (288Mb)
    - Separate 72-bit read and write ports
  - 350 MHz bus operation, DDR (double data rate)
    - Fast enough to be clocked at 312.50 MHz
    - Eliminates clock synchronization delays between memory and Ethernet clock
- Full support for embedded logic analyzers via JTAG interface
  - ChipScope and other third-party debug solutions:
    - InPA, Veridae, SpringSoft
- Status FPGA-controlled LEDs
  - Enough light to cause severe migraine headaches

Table 1: FPGA Resources

FPGA	Speed Grades (slowest to fastest)	LUT Size	FF's	Gate Estimate		Max I/O's	Multipliers (25x18)	Memory		
				Max (100% util) (1000's)	Practical (60% util) (1000's)			Blocks (18kbits)	Total (kbits)	Total (kbytes)
HX565T	-1,-2	6-input	708,000	6,797	4,078	720	864	1,824	32,832	4,104
HX380T	-1,-2,-3	6-input	478,000	4,589	2,753	720	864	1,536	27,648	3,456

### Block Diagram



## Description

### Overview

The **DNPCle\_10G\_HXT\_LL** is a PCIe-based FPGA board designed to minimize input to output processing latency on 10Gb Ethernet packets. The primary application is for ultra low latency, high throughput trading without CPU intervention. Every possible variable that affects input to output latency has been analyzed and minimized. Raw 10 GbE Ethernet packets can be analyzed and acted upon without interrupts or an operating system adding delay to the process. This configurable hardware computing platform has the ability to achieve the theoretical **minimum** Ethernet packet processing latency.

### The FPGA – Xilinx Virtex-6 HXT

We use a single FPGA from the HXT sub-family of Xilinx Virtex-6 in the FFG1923 package. This package supports 720 I/O with the majority utilized. Most are dedicated to a variety of off chip memory peripherals including QDR II+ for low-latency, high speed look-up and DDR3 for performance oriented bulk storage. The HXT FPGAs contain high-speed transceiver PHYs of two different types. GTX transceivers are capable of handling data rates of 150 MB/s to 6.5 Gb/s, making these useful for lower speed Ethernet and GEN1/GEN2 PCI Express. The GTH transceivers are tuned higher, 2.488 to 11 GB/s, making them applicable to 10 gigabit Ethernet (10 GbE). Eight of the GTX transceivers are used for GEN2-capable PCIe. Three of the GTH transceivers are connected to 10 GbE SFP+ sockets and a fourth to a CX4 socket. Another 8 GTX transceivers are connected to an expansion connector, allowing for peripheral expansion but most applicable to in-chassis, board to board data daisy chaining.

Two possible FPGAs can be stuffed: HX380T or the HX565T. The HX380T comes in three speeds grades, with -3 being the fastest. The larger HX565T is limited to the -2 speed grade. This means the smaller device can be clocked at a higher frequency at the cost of slightly fewer FPGA logic resources. Table 1 depicts the resources of the two FPGAs with the Xilinx marketing exaggerations removed. These are both large FPGAs. The HX565T is capable of handling >4M ASIC gates of logic and is among the largest of the FPGAs shipping from any vendor in 2011. Features of the Virtex-6 HXT FPGAs include the efficient, dual-register 6-input look-up table (LUT) logic, 18 Kb (2 x 9 Kb) block RAMs, and second generation DSP48E1 slices (includes 25 x 18 multipliers). Floating point functions can be implemented using these DSP slices.

To give you an idea as to how large these FPGAs are, Xilinx has embedded processor IP called [MicroBlaze](#). This processor is implemented in FPGA logic gates. Fifty (50!) or more of these MicroBlaze processors can be stuffed into an HXT565T with room to spare. Somewhat fewer if you incorporate IEEE 754 floating point.

### Three Channels of 10 GbE, One channel of Infiniband

The HXT FPGAs have transceivers capable of 10 GbE. The physical interface is handled using SFP+ modules. This allows you to bypass a MAC if necessary and process raw Ethernet packets. The **DNPCle\_10G\_HXT\_LL** has 3, 10 GbE channels and can support 10GBASET-ER, 10GBASET-SR, 10GBASET-KR. Another 10 GbE connector is CX4, enabling low latency network connections via Infiniband.

### QDR II+ SSRAM

We use 4 individual quad data rate static RAMs (QDR II+ SSRAM) in the 4M x 36 size, configured in three independent blocks. Two of the memories are configured as 4M x36, and two more are combined as 4M x 72. This style of memory has separate input and output data paths, enabling maximum read/write data bandwidth with minimum latency. These three separate memories units can be controlled individually, but any two (2M x 72) or all three can be combined and treated as a single memory. The maximum tested frequency of this memory is 400 MHz. To minimize processing latency, we suspect it will be best to clock these QDR II+ SRAMs at 312.50 MHz, exactly twice the internal Ethernet controller frequency of 156.25 MHz. The HXT FPGAs are capable of generating internal 2x clocks that are phase synchronous, eliminating the latencies associated with the tricky re-synchronization of data moving between different clock frequencies. The internal controller can be optimized in any way you choose. We, of course, provide several verilog examples for no charge that you are welcome to use. All functions of the QDR II+ SSRAM can be exploited, including concurrent read and write operations and four-tick bursts. The only real limitation is the amount of time and effort spent in customizing the individual memory controllers.

**DDR3 – A large amount of local, bulk memory**

A single PC3-8500 DDR3 DIMM socket enables up to 16GB of memory for bulk storage and lookup. Assuming a 16GB DIMM, the memory configuration is 2048M x 72. Using a -2 or -3 speed grade FPGA, this interface is tested at the maximum FPGA I/O frequency: 533 MHz (1066 Mb/s with DDR). You are welcome to use this memory as 64-bits with 8 bits of error correction (ECC), or as a 72-bit memory without correction.

To minimize data synchronization across clock boundaries, it probably makes sense to clock this DDR3 interface at a 3x multiple of the base Ethernet frequency of 156.25 MHz, which is 468.75 MHz. A 3x phase synchronous clock can be easily generated internal to the FPGA, allowing zero latency synchronous data transfers between the Ethernet packet receiving logic and the DDR3 memory controller. The DDR3 controller can be optimized in any way you choose. We, of course, provide several verilog examples for no charge that you are welcome to use. All functions of the DDR3 DRAM can be exploited and optimized. Up to 8 banks can be open at once. Timing variables such as CAS latency and precharge can be tailored to the minimum given your operating frequency and the timing specification of the exact DDR3 memory utilized. As with the QDRII+ SRAM, the only real limitation is the amount of time and effort spent customizing the DDR3 memory controller to your needs.

**PCIe – Customizable 8-lane, GEN2 PCI Express**

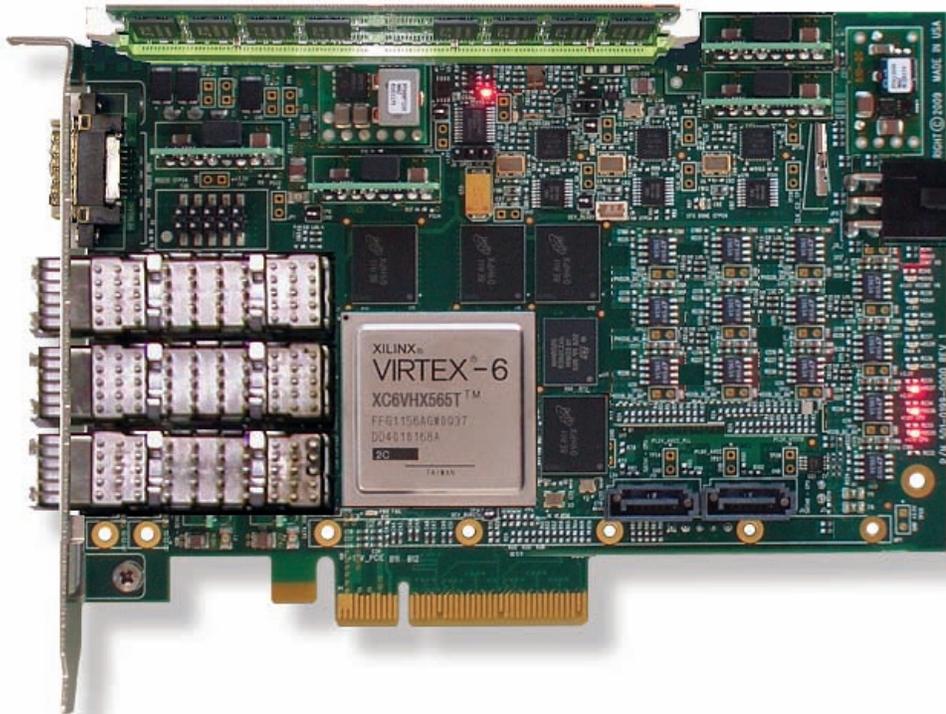
PCIe is connected directly to the FPGA via 8-lanes of GTX transceivers. The interface is fully GEN2 capable. We ship PCIe IP that is a full function, fixed, 8-lane master/target. To gain access to the PCIe interface, this IP must be integrated with your application. We can help configure this IP to your needs, including BAR sizes. Additionally we can optionally add or subtract DMA engines, scratchpad memories, interrupts, and other host-related functions to maximize the performance, while utilizing the minimum FPGA resources. Drivers and ‘C’ source for several operating systems are included no charge. Partial reconfiguration of the FPGA is supported via the PCIe interface.

**Board to Board Daisy Chaining and Expansion**

These boards can be stacked in a PCIe system utilizing the GTX Expansion Header. We connect four lanes of the GTX transceivers to a high speed connector. This enables high board to board communication at the rate of 10 Gb/s on each of the four lanes and bidirectional.

**How Everything Works ....**

With direct data feeds such as NASDAQ ITCH and OUCH, the **DNPCIE\_10G\_HXT\_LL** contains all of the basic functions required to minimize the amount of time it takes to receive Ethernet packets, process them, and respond deterministically. The MAC, operating system et al, can be bypassed. There are no interrupts. No operating system. Not a single clock cycle is wasted here, enabling a near theoretical minimum in-to-out response time. For algorithms requiring processing, FPGA resources can be hard coded to perform the task. This includes real-time Monte Carlo analysis and floating point, all operating 1000's of times faster than possible in a processor-based approach.



---

For technical applications and sales support, call 858.454.3419

**DINI**  
Group

7469 Draper Ave.  
La Jolla, CA 92037-5026  
Phone: 858.454.3419  
Fax: 858.454.1728  
E-Mail: [sales@dinigroup.com](mailto:sales@dinigroup.com)  
Web: <http://www.dinigroup.com>

*The DINI Group reserves the right to make changes to the product(s) or information contained herein without notice. No liability is assumed as a result of their use or application. No rights under any patent accompany the sale of any such product(s) or information.*