MANCHESTER 1824

The University of Manchester

DAQ Electronics

Marc Kelly
DAQ Group, Particle Physics, University of Manchester.

MANCHESTER 1824

The University of Manchester

Some Goals

- Exploring the ability of modern FPGAs to drive high speed Ethernet as a means of data accusation in future detector readout.
- Gain understanding of the advantages and drawbacks of using such a system.
- Produce a working testbed for running various tests at 1GBit rates, and a future expansion into 10GBit systems.
- Provide useful firmware and additional modules that can be used by other areas of the DAQ development effort.

MANCHESTER 1824

The University of Manchester

Ethernet In FPGA

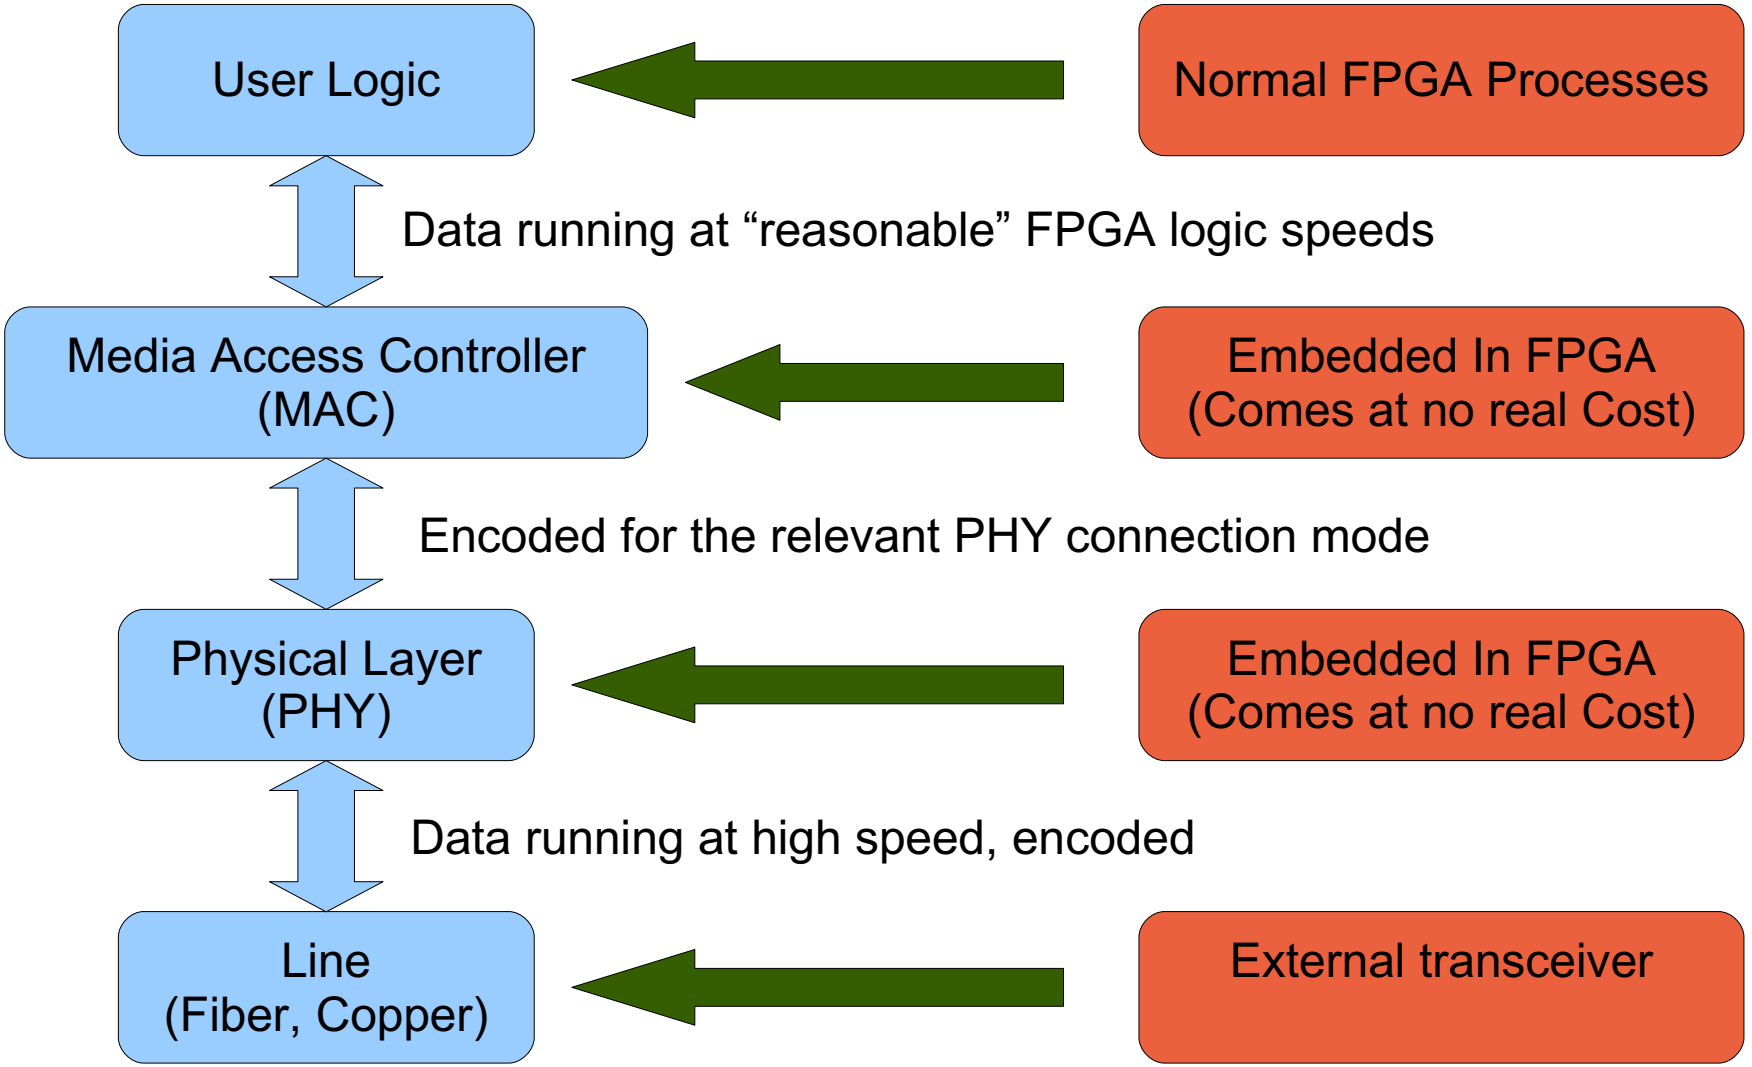

MANCHESTER 1824

The University of Manchester

Xilinx Virtex-4 FX

- Has 2 embedded Gigabit capable MAC layers. These exist as hardware primitives, and come at no real user logic cost.
- PHY layer can be provided using the embedded high speed “RocketIO” serial encoder/decoders (SERDES).
- Contains all the required logic to perform various protocols for PHY connections, MII GMII SGMII or 1000BASE-X. User logic does not have to deal with the data encoding, the MAC layer does it.
- RocketIO can directly provide the 1000BASE-X output for connection to external Transceivers. This includes all the 8B/10B encoding and clock recovery.
- This allows the use of either fibre or copper connections with the right tranceivers. Currently we are using copper.

MANCHESTER 1824

The University of Manchester

Pics 1

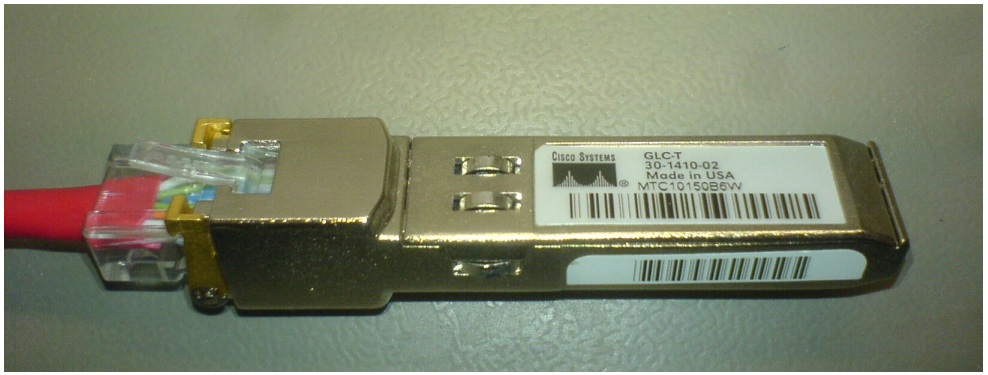

A 1000Base-X to 1000Base-T transceiver

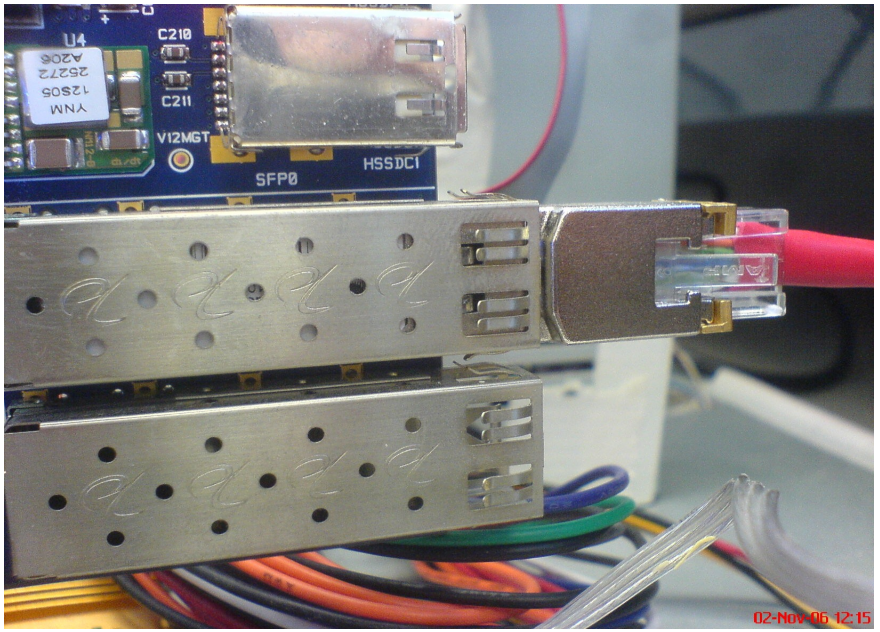

The transceiver inserted into the cage on the FPGA board.

MANCHESTER 1824

The University of Manchester

Pics 2

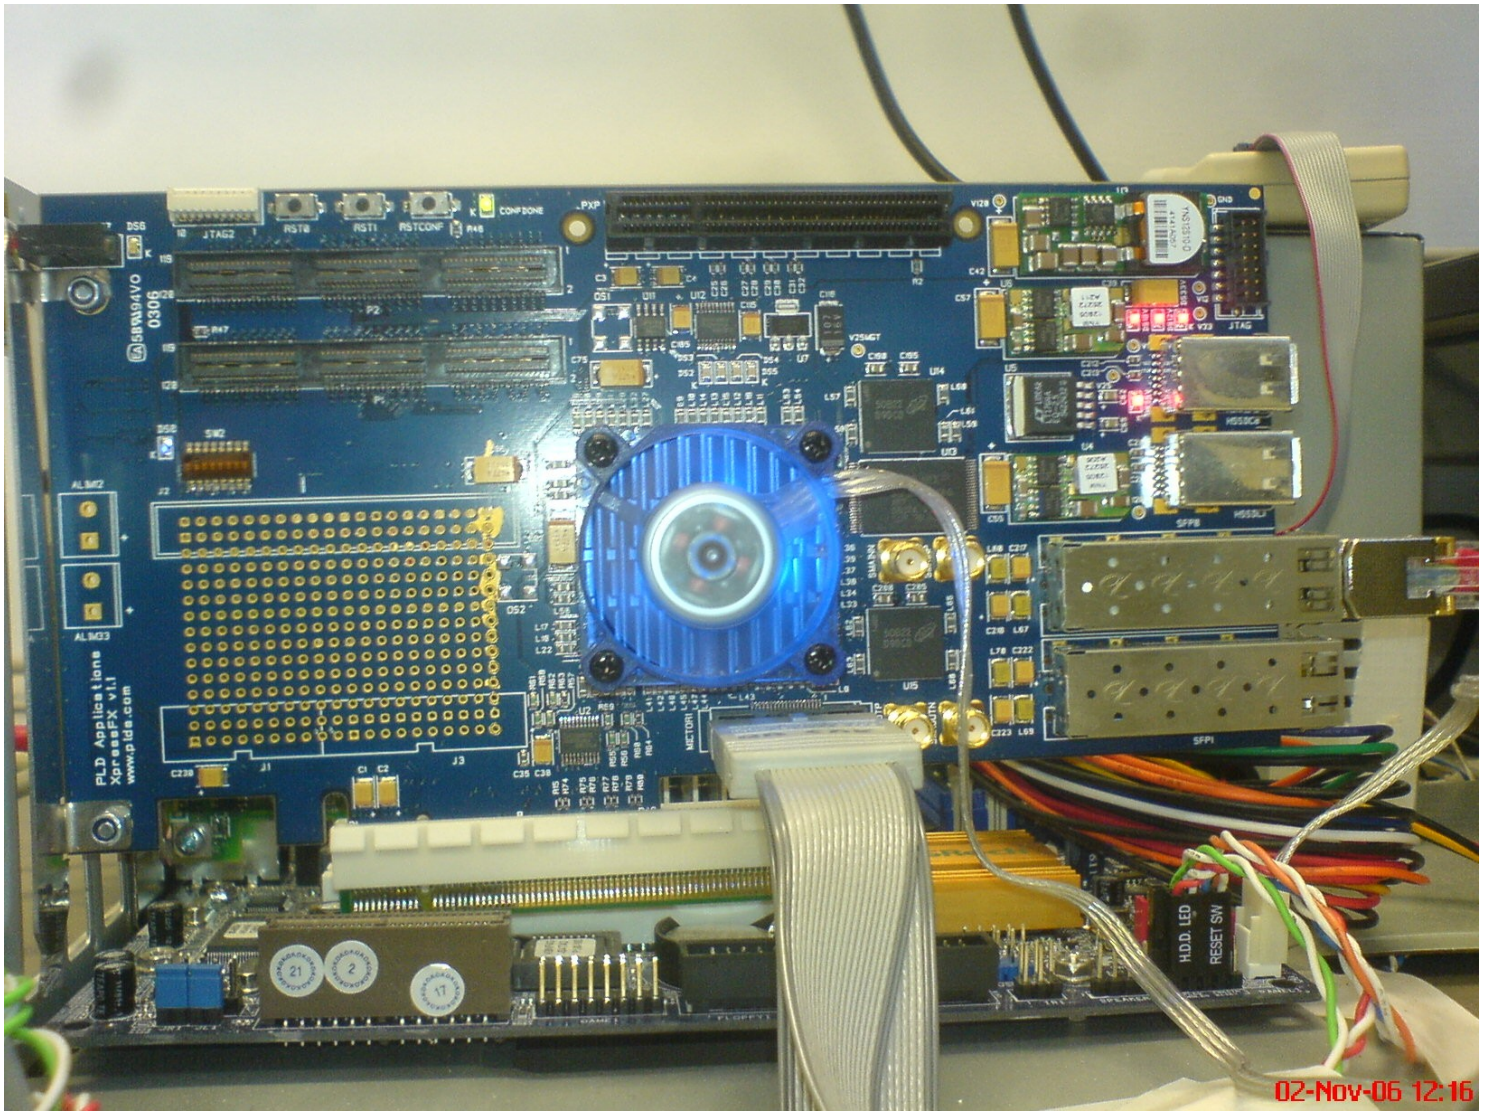

Full FPGA Board Inserted into a host computer (This only powers the board)

Virtex 4 under cool blue heatsink.

MANCHESTER 1824

The University of Manchester

Test Firmware Design 1

- A reusable module containing the MAC, PHY and associated control logic.
- A generic Packet Generator and FIFO system to allow packets to be sent and received from the MAC layer. This is interchangeable, and allows us to change the functions in an easy manner. For instance a simple “Ping” style loopback, can be used in place of the packet generator.
- A simple Embedded CPU (Based on Motorola 68HC11) is used to provide a method of performing configuration and control. This also provides some basic Ethernet communications to the outside world.
- Various parts of the firmware can be reused by the DAQ receiver project.

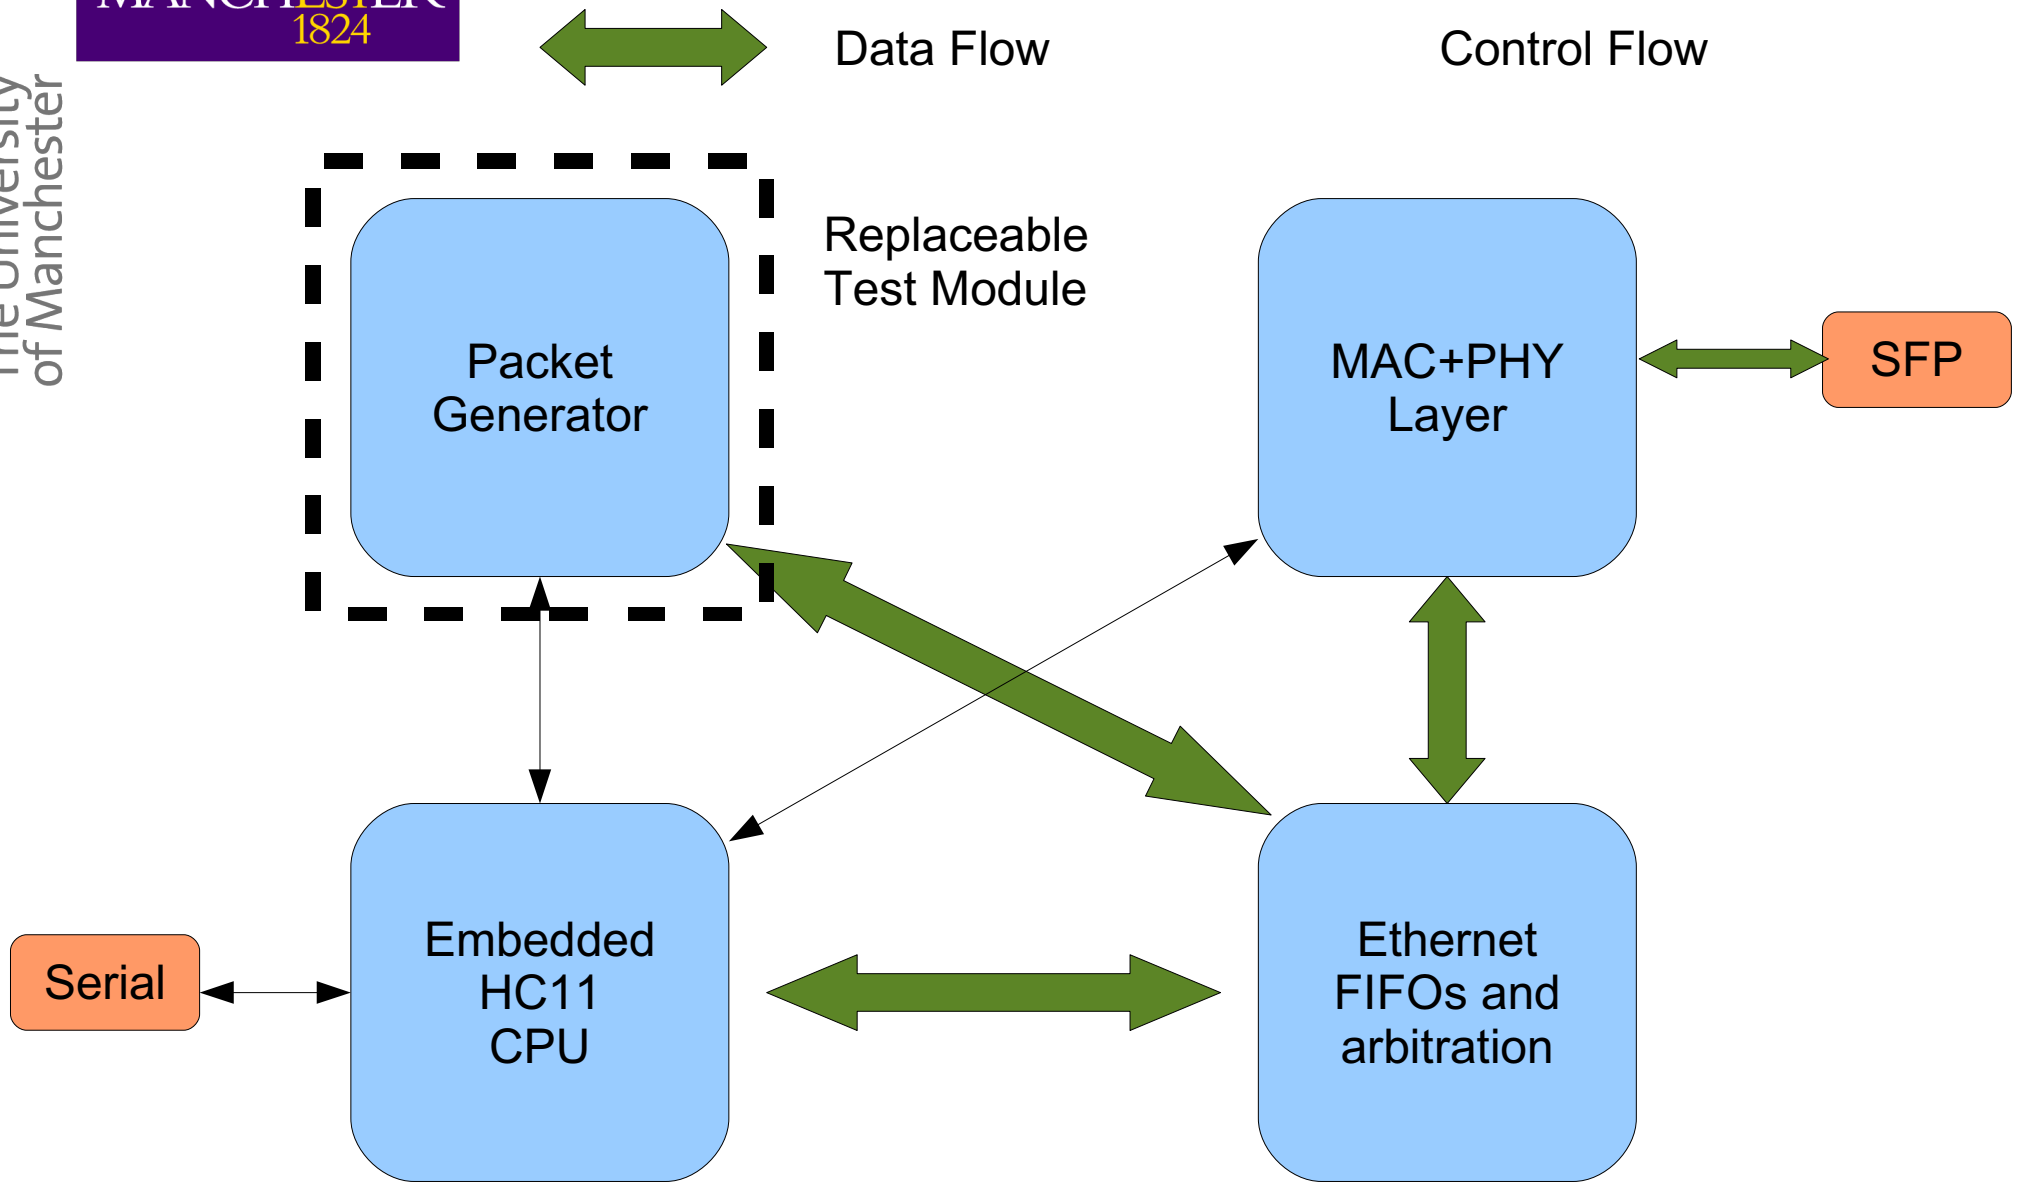
Data Flow
Control Flow
Replaceable Test Module
Packet Generator
MAC+PHY Layer
SFP
Serial
Embedded HC11 CPU
Ethernet FIFOs and arbitration

MANCHESTER 1824

The University of Manchester

Test Firmware Design 2

- The current packet generator module allows up to 16 destination MAC addresses to be used, and allows packet count/size/delay to be controlled.
- Pseudo random data can be delivered in the packets from a Linear Feedback Shift Register, The seed can be configured. This can be used to examine the number of bit errors since the expected sequence can be calculated at the receiver end.
- A simple IP style checksum is calculated via a simple offload engine and appended to each packet to allow bad packet testing. This is a test module for a future requirement if we move to UDP packets instead of RAW Ethernet.

MANCHESTER 1824

The University of Manchester

Test firmware Design 3

- By using an embedded CPU to control the packet generator it allows us to try several different methods without having to totally rebuild the firmware for each.
- This is done by recompiling the code to run on the HC11, and then using Xilinx “Deep Magic” patch the existing placed and routed bitstream to incorporate the new software.
- Very useful when testing such things as auto-negotiation before moving onto large scale packet testing.
- As the HC11 also has access to the MAC's data stream, this also allows us to do other things such as add diagnostics which can output via the Ethernet route, or via an embedded UART to a serial console.

MANCHESTER 1824

The University of Manchester

Test Setup

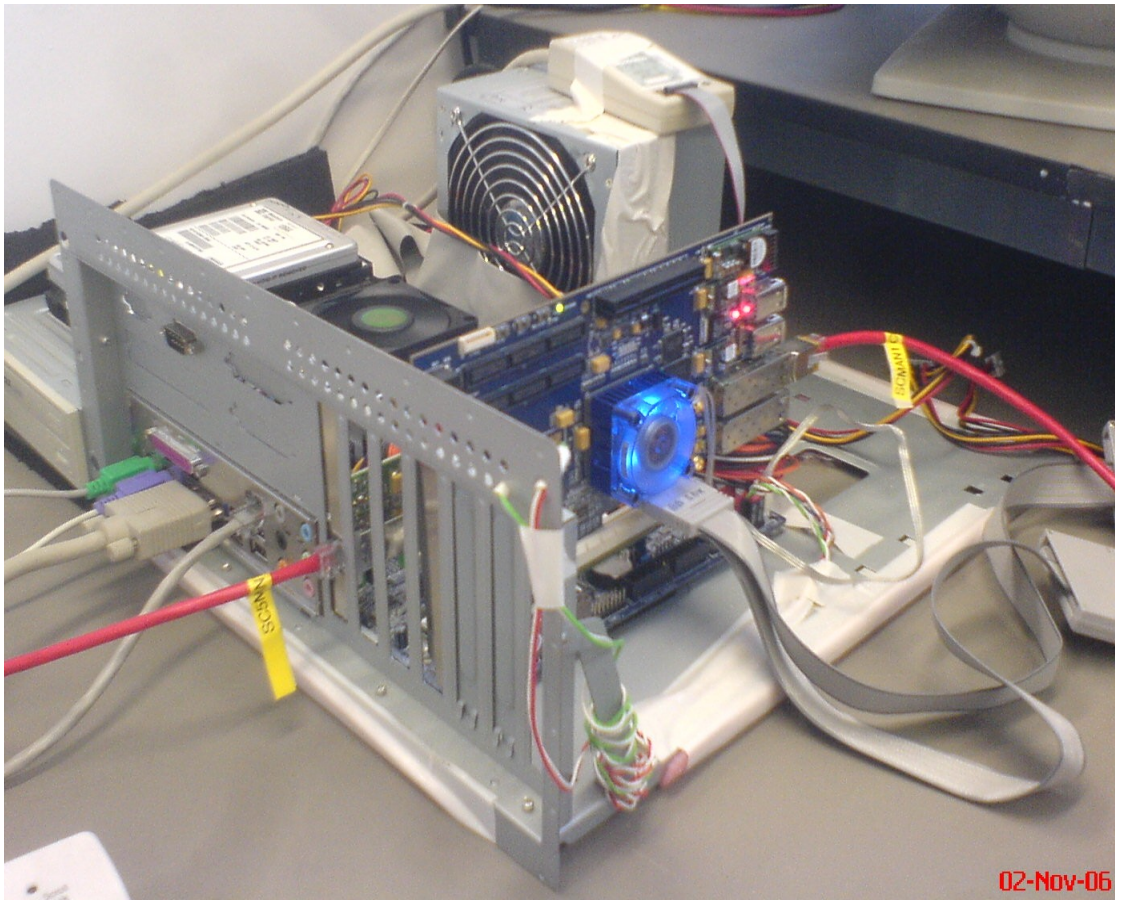

The test stand used for initial development. The FPGA card is installed in the machine, and there is a PCIexpress Gigabit Ethernet card installed.
The Ethernet card is known to be able to perform at line speed from previous testing done by RHJ.

MANCHESTER 1824

The University of Manchester

Other network cards

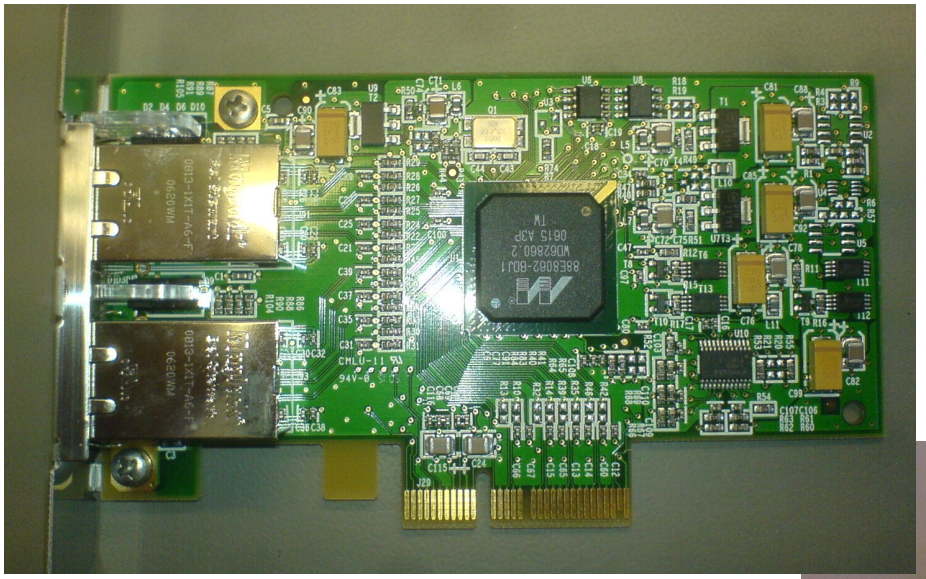

Syskonnect PCIExpress Card, has duel Ethernet capability on a 4x connector

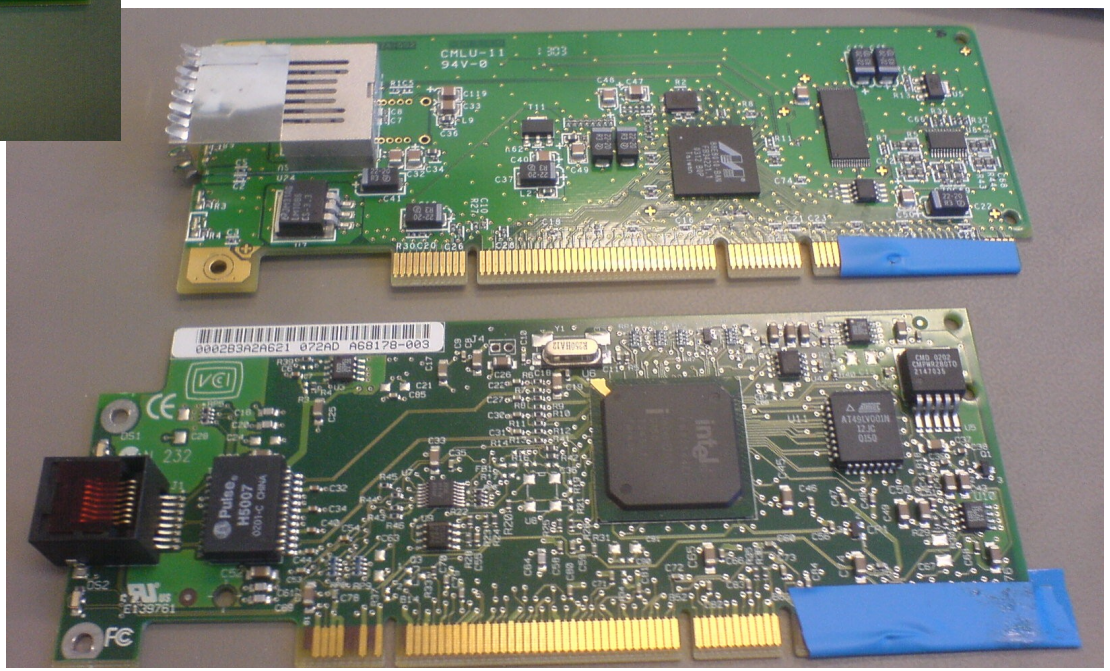

Intel, Syskonnect PCI-X cards. Run in PCI motherboards they are bandwidth limited. Also have Netgear card (not shown)

MANCHESTER 1824

The University of Manchester

Testing 1

- A selection of FPGA board to single PC, back to back and via a switch. With various packet sizes, delays and network cards. This will allow us to test the performance of the receiver system when dealing with RAW Ethernet. And to prove that an FPGA can cope with sending such data rates, and still be able to perform other tasks.
- FPGA board to multiple PC receivers. To demonstrate we can distribute data to many destinations at line rate.
- Utilize the second MAC layer in the FPGA, doubling the data output rate.
- Produce, or acquire a 10GBit expansion board, and 10GBit MAC firmware to allow testing of machines that much higher data rates.

Testing 2

- Investigate extra functions, such as adding data redundancy in the transmitted stream to cope with bit errors and lost packets. A simple method would be something equivalent to the method used by RAID 5, where we can suffer a single packet loose in a stream of length N. You do however loose bandwidth as you are sending redundant data.
- Examine the prospect of data compression (gzip, RLE etc.) to reduce the data rate required. This requires the use of the PowerPC's and DDR memory that the board provides.
- When running with a 10GBit output have a system where by we take multiple 1Gig inputs, aggregate them and then fan them out at 10Gig via a switch.

Status

- Firmware mostly written for testing as a transmitter.
- Now have a good knowledge of the very fickle nature of the RocketIO system and all its horrible features.
- Embedded HC11 Assembler/C code to run as a transmitter is in progress.
- Software for Linux receiver stations in progress.
- Initial system auto-negotiation to various Linux network cards has been demonstrated and proven.
- Initial packet TX/RX using a pre-built Xilinx “loop back” module has been demonstrated
- High speed packet generation has been Show, although no testing on these packets was performed. This was done by generating random packets and monitoring the NIC using “WireShark”, aka Ethereal.