MANCHESTER 1824

The University of Manchester

Thermal, Mechanical & Networking Issues

Dave Bailey
University of Manchester

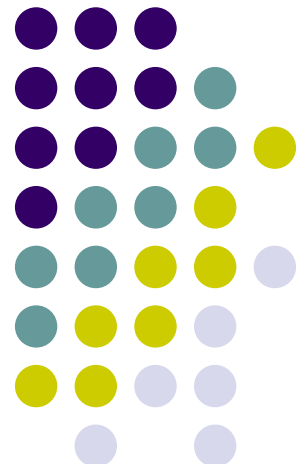

MANCHESTER 1824

The University of Manchester

Outline

- Thermal Simulations
- Mechanical Issues
  - Long-term studies of conductive glue
  - Automated assembly
- Few words on 10 Gig Ethernet investigations

MANCHESTER 1824

The University of Manchester

CALICE Cooling

Steve Snow

5/5/06

A CALICE module will dissipate at least 300 W so it will need some active cooling.

The obvious place to cool it is at this end. One could cool either the slabs or the alveolar structure or both. Disadvantage is that this end is already busy with slab readout.

Alternatively one might cool this end. Disadvantage is dead material within calorimeter active volume and not possible to cool the slabs directly.

Cooling of this face should not be ruled out. Disadvantage is relatively poor conductivity of module in the perpendicular direction.

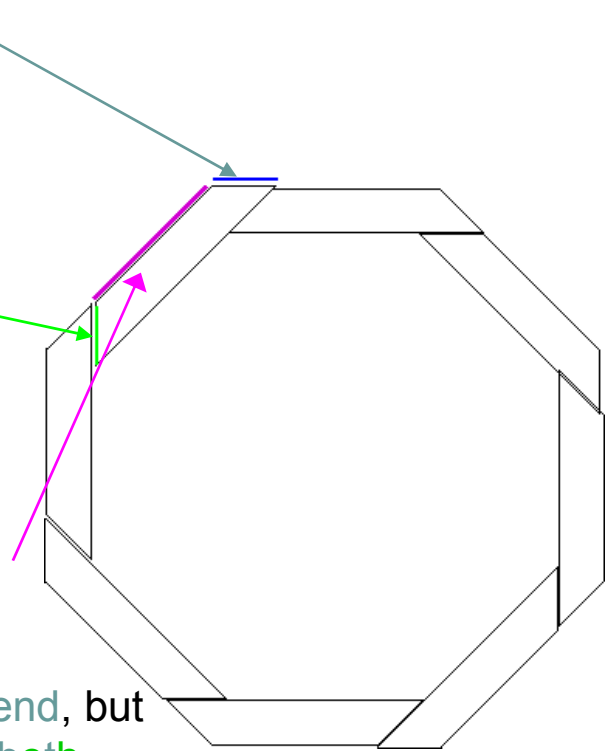

The rest of this talk will assume cooling at the most obvious end, but results can easily be extended to cooling at the other end or both.

Overview

By far the best thermal conductors for carrying heat along the slab direction are the almost continuous plates of tungsten. If the heat generated in the front-end chips can easily reach the tungsten then nearly all of it will go this way and the temperature gradients are simple to calculate. This talk will show that the heat can indeed easily reach the tungsten.

First, some order of magnitude estimates:

By far the worst thermal conductor in the CALICE module is the air. Air layers must exist between the slab and alveolar structure for clearance when the slab is inserted or removed. A reasonable estimate for the clearance is 0.3mm. When the slab is in place, the air gap will be zero at a few contact points, but on average it will be half the clearance.

The thermal resistance of this air gap, over a 6cm x 6cm unit cell area is:

$$R_{air} = 0.15 \text{ mm} / (0.024 \text{ (W/m/K)} \times 0.06 \text{ m} \times 0.06 \text{ m}) = 1.7 \text{ (K/W)}$$

Compare this with the thermal resistance of the tungsten that joins one cell to the next along the slab direction:

$$R_W = 0.06 \text{ m} / (177 \text{ (W/m/K)} \times 2.1\text{mm} \times 0.06 \text{ m}) = 2.7 \text{ (K/W)}$$

So the geometrical factor approximately balances the conductivity factor. When you consider that the heat must cross one R_{air} and up to 26 R_W resistances to reach the coolant, it is clear that getting into the tungsten is the easy bit.

Estimates continued

We can continue this order-of-magnitude estimate by

- negelcting R_{air}
- assuming that all the heat flows in the tungsten so the thermal resistance of the cell, R_{cell} , is equal to R_W
- assuming there is a constant heat input into each cell, P_{cell} = 0.015 W, corresponding to 0.1 mW /channel and 5x5 mm^2 pads.

In this situation, the temperature difference along a row of N cells that are cooled at one end is

$$\Delta T = 0.5\, P_{cell}\, R_{cell}\, N^2$$

Assuming a module is 26 cells long, the temperature difference from one end to the other will be around 13.7 C. If it was cooled at both ends the temperature of the middle relative to the ends would be around 3.4 C.

Finite element simulation

We now go on to simulate the heat flow in detail, using FlexPDE.

Our model of the cell geometry is based on information from Marc Anduze and Catherine Clerc.

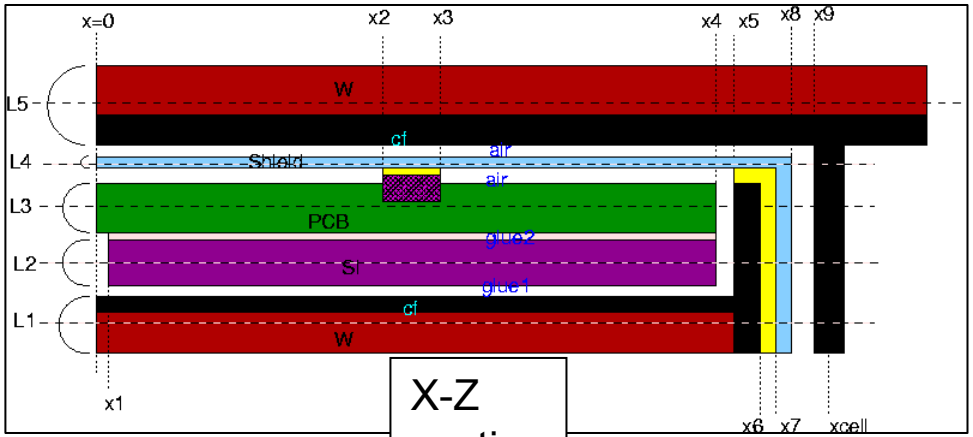

X-Z section

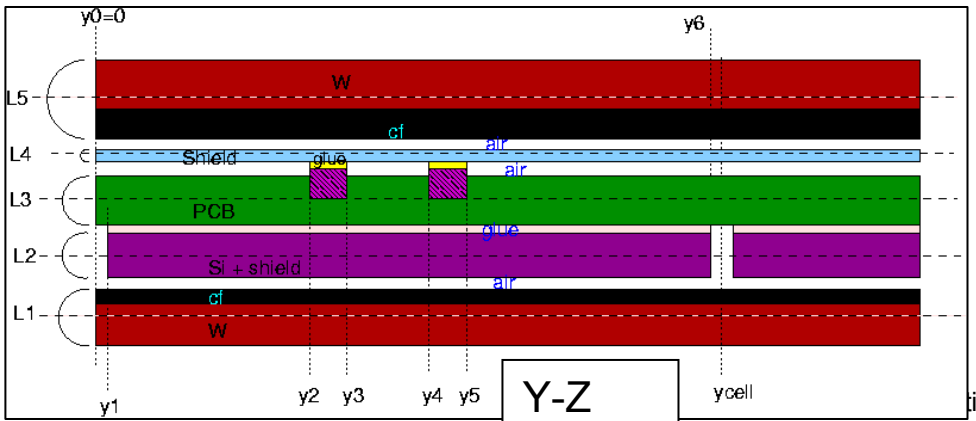

Y-Z section

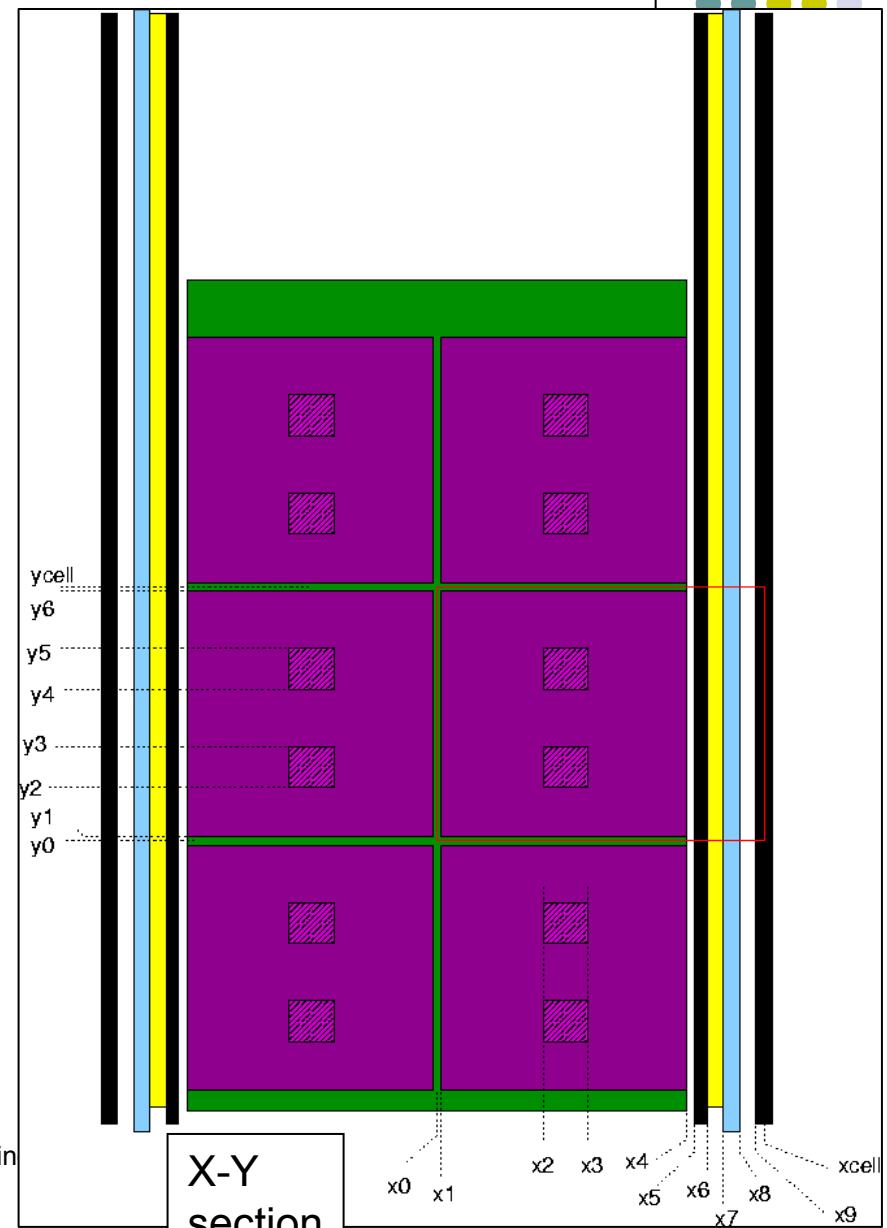

X-Y section

MANCHESTER 1824

The University of Manchester

Dimensions and Conductivities

For completeness, here are all the dimensions and conductivities, taken directly from the simulation code.

Note that the units are Watt, Kelvin and millimeter.

```
{ Dimensions in X,Y.}
      w_Sigap   = 0.15 { width of gap between detectors on same PCB}
      w_chip    = 15.0 { width of chip }
      w_Si      = 62.0 { width of detector }
      w_edge    = 0.15 { width of air gap between detector edge and CF H edge}
      w_CFHU    = 0.3  { width of the uprights of the CF H frame }
      w_glue_Al = 0.1  { width of glue layer holding Al shield }
      w_Al      = 0.1  { width of the aluminium shielding }
      w_clear   = 0.15 { width of clearace gap between slab and the alveolus wall }
      w_CFalve  = 0.6  { width of the alveolus upright CF walls }
      w_PCBgap  = 0.2  { width of gap between PCBs }
 { Dimensions in Z. }
      th_W_slab  = 2.1  { thickness of tungsten in slab }
      th_W_alve  = 2.1  { thickness of tungsten in alveolus wall }
      th_CF_HH   = 0.3  { thickness of CF in horizontals of H frame }
      th_CF_alve = 0.3  { thickness of CF in horizontals of alveolus }
      th_Si      = 0.555{ thickness of detector, including ground foil }
      th_PCB     = 0.7  { thickness of PCB }
      th_Chip    = 0.8  { thickness of chip }
      th_Cutaway = 0.1  { thickness of PCB in cutaway under chip }
      th_Chipgap = 0.1  { thickness of chip to shielding gap }
      th_g_SiH   = 0.05 { thickness of glue dots joining detector to H frame}
      th_g_SiPCB = 0.05 { thickness of glue dots joining Si to PCB }
      th_g_shield= 0.1  { thickness of glue joining shield to H frame }
      th_a_PCBSh = 0.3  { thickness of air gap between PCB and Al shielding }
      th_shield  = 0.1  { thickness of shielding }
      th_a_clear = 0.15 { thickness of air for clearance between slab and alveolus }
 { Fraction of area covered by glue }
      cov_g_SiH  = 0.02  { coverage of glue dots joining detector to H frame }
      cov_g_SiPCB= 0.2   { coverage of glue dots joining detector to PCB }
 { Thermal onductivities }
      C_W      =   0.177    { conductivity of tungsten    W/ (mm*K)  }
      C_Si     =   0.168    { silicon }
      C_Al     =   0.180    { aluminium }
      C_Cu     =   0.385    { copper }
      C_N      =   0.000024 { nitrogen or air }
      C_CF_in  =   0.00324  { carbon fibre in-plane }
      C_CF_out =   0.00061  { carbon fibre out-of-plane }
      C_shield =   C_Al
      C_g       = 0.0002     { a typical glue }
      C_gAg     = 0.005      { an electrical conductive glue }
      C_g_SiH   = cov_g_SiH*C_gAg + (1-cov_g_SiH)*C_N          { glue  Si to H }
      C_g_SiPCB = cov_g_SiPCB*C_gAg + (1-cov_g_SiPCB)*C_N      { glue Si to PCB }
      C_g_HSh   = C_g        { glue Al foil to H }
      C_PCB_in  = 0.025      { PCB in-plane }
      C_PCB_out = 0.0003     { PCB out-of-plane }
      C_chipgap = C_g
```

Simulation features

We use a laminar approximation, in which temperature is a function of X and Y but not Z within each layer. The variation of temperature with Z is described by using five separate but coupled temperature fields. The labelling of the layers is:

Layer 1 is the tungsten of the slab and its carbon fibre
Layer 2 is the silicon sensors including the ground foil
Layer 3 is the PCB and includes the readout chips
Layer 4 is the shield that encloses the slab
Layer 5 is the tungsten of the alveole and its carbon fibre

In order to complete the model one needs a heat source and sink:

In some versions of the model we connect all layers to a heat source on one side of the cell and to a sink on the other side. This simply measures the thermal resistance of the whole cell, assuming that heat can easily get to any layer.

In other versions we connect one or more layers to a heat sink on one side of the cell and make the front-end chips the source. This is used to check the assumption above and quantify its accuracy.

Thermal resistance

The simulation predicts that the thermal resistance of the whole cell is

$$R_{cell} = 2.03 \text{ K/W}$$

with almost 80% of the heat flowing in the tungsten and most of the remainder in layers 2 + 3 (silicon + PCB).

Temperature map

Here is a typical temperature map from a simulation of three adjacent cells, cooled at the top end.

In this simulation heat is input to the cells at the rate of 1 W per cell, whereas the power expected in CALICE is 0.015 W/cell, so temperatures should be scaled accordingly.

Heat sink.

Connection from cells to heat sink. Can be turned on or off. In this run of the simulation all connections were on.

Cell 1

Cell 2

Cell 3

FE chips position

average chip temperatures reported

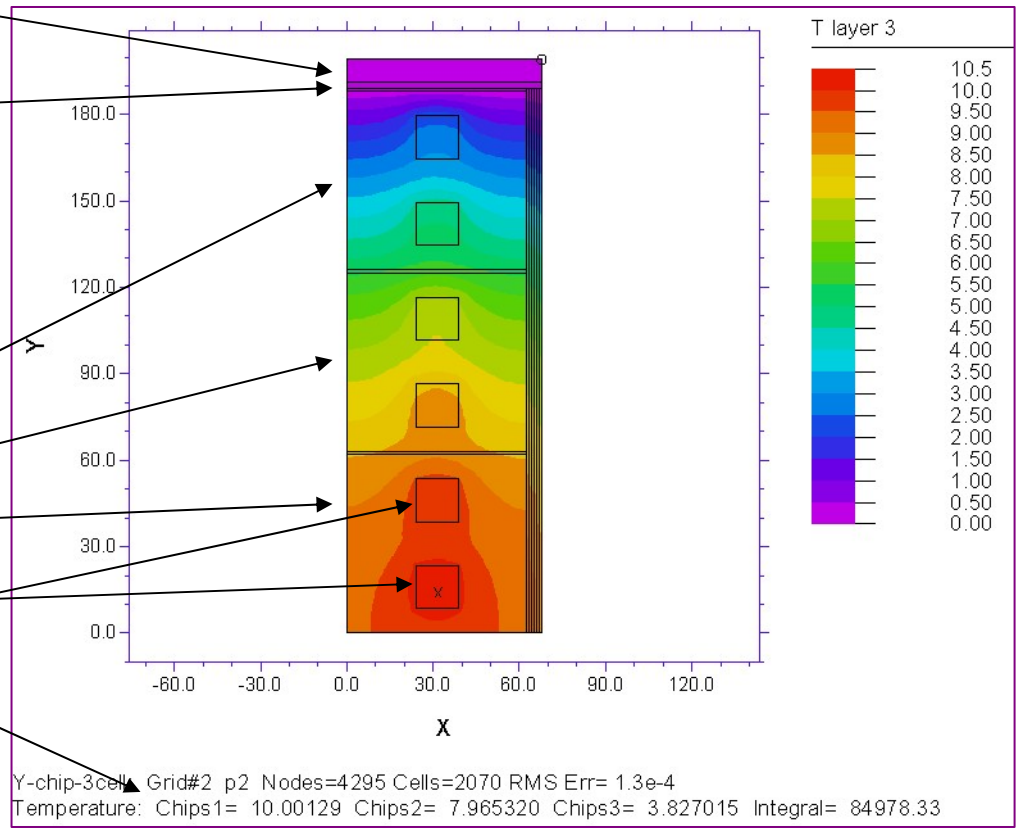

MANCHESTER 1824

The University of Manchester

Results

Layers connected to heat sink	Temperatures relative to heat sink			Percentge of heat flowing in layers				
	Chip 1	Chips 2	Chips 3	L1	L2	L3	L4	L5
All	3.83	7.96	10.00	52.5	0.6	6.9	4.8	35.3
L1 and L5	4.11	8.14	10.17	59.8	0	0	0	40.2
L1 only	6.51	11.18	13.34	100	0	0	0	0
L5 only	11.22	14.29	16.12	0	0	0	0	100
Model 2 K/W/cell	3	7	9					

First 4 rows of this table are from FlexPDE simulation with different layers connected to the heat sink. Last row is a simple model in which the chips positions are taken as ½, 1 ½ and 2 ½ cells from the heat sink and the cell resistance is taken as 2 K/W.

The plot shows that the difference between the simple model and the more detailed simulation is a temperature offset, which is smaller if the main conductive layers are connected to the heat sink. The offset is due to the thermal resistance between the FE chips and the conductive layers.

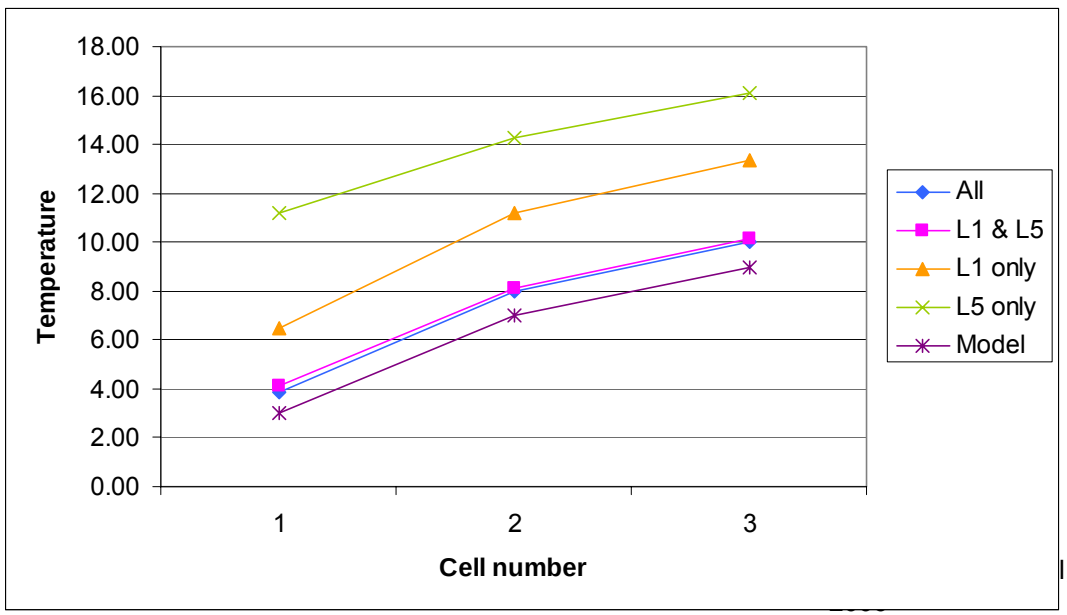

To get from here to the full CALICE module one must:

1. multiply T by 0.015
2. extrapolate from 3 to 26 cells.

This extrapolation sounds extreme but is in fact very safe, because as one goes to at larger number of cells the heat has more chance to get into the most conductive layers.

Thermal Simulation Conclusions

A detailed simulation shows that the thermal resistance of a CALICE cell in the direction along the slab is R_{cell} =2.03 K/W. This is slightly lower than our initial estimate of 2.7 K/W because some heat flows in materials other than the tungsten.

There will be a temperature offset difference between the chips and the cooled layer(s) of order 0.1 to 0.2 K, depending on which layers are connected to the heat sink.

In this situation, the temperature difference along a row of N cells that are cooled at one end is

$$\Delta T = 0.5\, P_{cell}\, R_{cell}\, N^2$$

Assuming a module is 26 cells long, the temperature difference from one end to the other will be around 10.3 C. If it was cooled at both ends the temperature of the middle relative to the ends would be around 2.6 C.

There is very little that can be done to either improve these temperature differences or make them worse, because they just depend on the dimensions and conductivity of the tungsten and the non-negligible conductivity of air.

MANCHESTER 1824

The University of Manchester

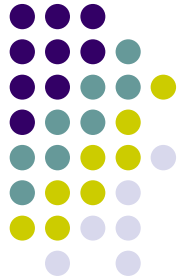

Mechanical Issues

- Glue Studies
  - Conductive epoxy extensively used in CMS
    - CMS Tracker modules
    - Conducting Glue used for bias connection to backplane
  - Issues have arisen during assembly

CMS Tracker Problems

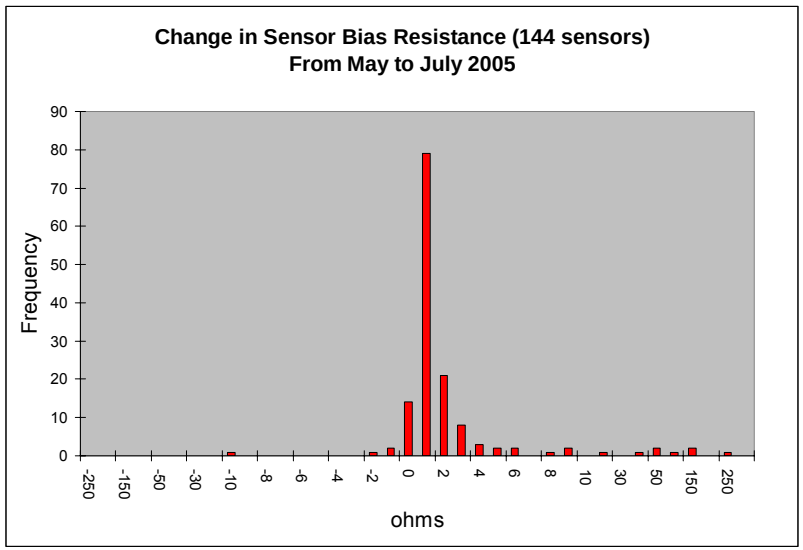

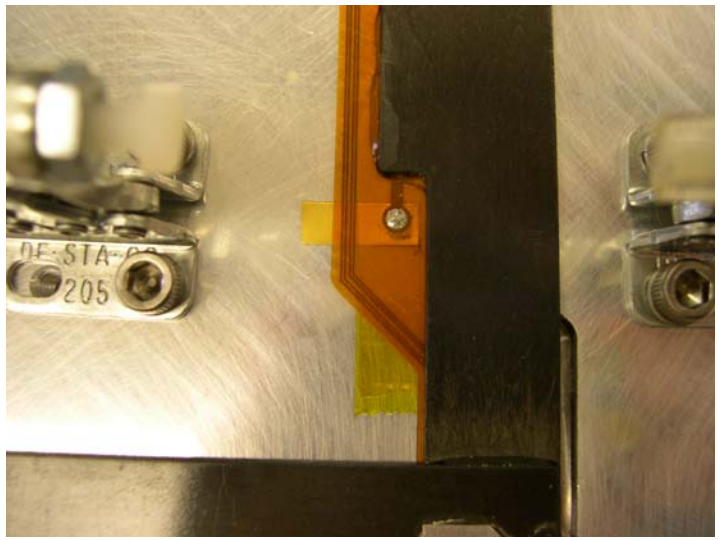

Tony Affolder, UCSB

- Of 51 TOB modules built, found 29 sensors with 1K to over 40MΩ higher than normal resistance in the bias connection.
  - Added strictly defined mixing procedures for the Tra-Duct 2902.
  - Applied epoxy immediately before sensor placement.
  - Switched to using Epotek 129-4 silver epoxy.
  - Evidence problem is glue to sensor interface, not glue to gold pad
- No effect – eventually wire bonded all connections

Glue Test Setup

- Programmable Environmental chamber
- Resistance measurements
  - Keithley 2000 DVM
  - Keithley 236 SMU
- Labview control system

Resistance Snake Tests

Silicon / al

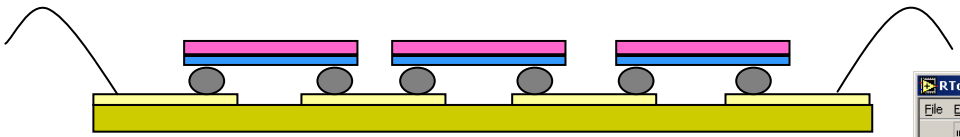

Pc board / gold

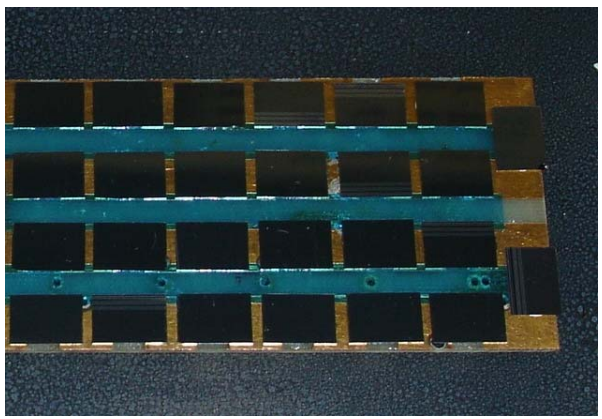

50 dots continuous monitoring

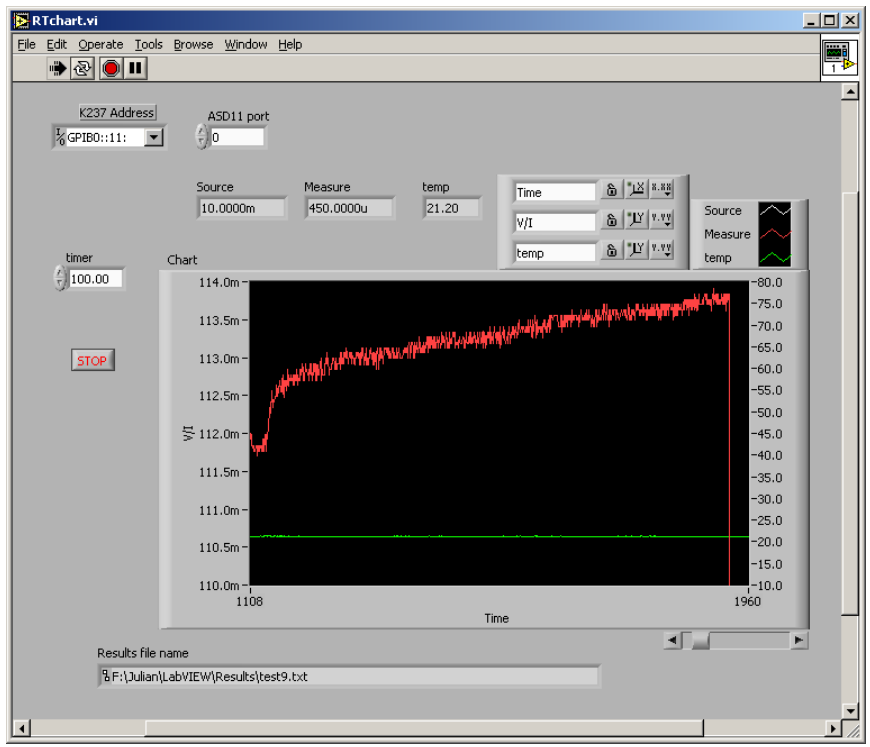

Temperature Cycling

- 20-70 1º/ min
- Upwards resistance rise
  - Consistent with silver temperature coefficient
- On contraction
  - Interfaces become loose?
  - Still under investigation

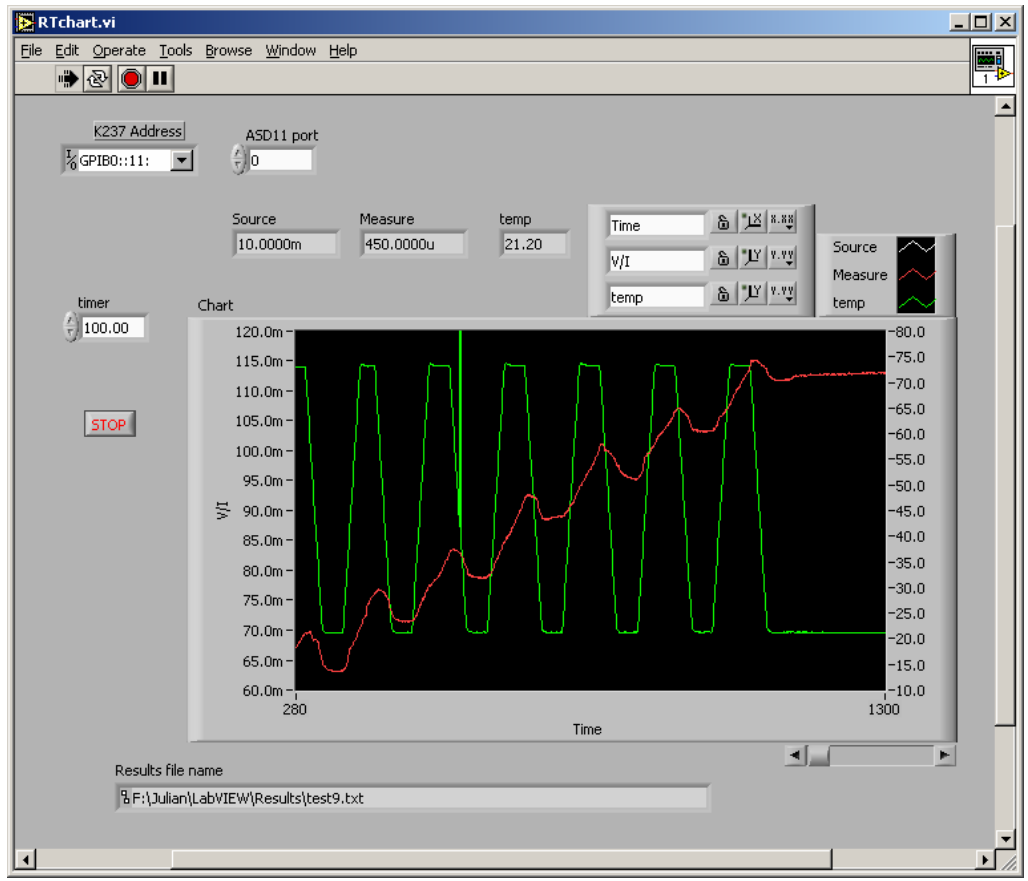

Further Steps

Silicon Wafer (when "spare" available)

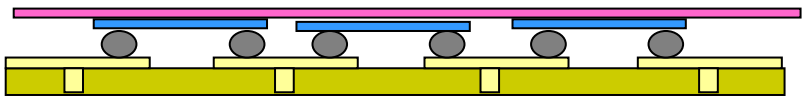

Double sided Pc
Measure lots of pairs after cycling

Would like to investigate availability of wafer samples:

Duff wafers, Mechanicals, Test structures at edges can be used instead of fully working silicon

Real stresses – info from Steve on temperature distributions

Standard glue tests (85°/85% humidity) on different conducting epoxies

Automated Assembly

- Start thinking about on pick and place techniques for assembling detector planks
  - Don't reinvent the wheel
  - Look at/ exploit existing CMS/Atlas experience
- Generalised concept of a gantry with interchangeable tools
  - Pattern recognition software for placement

Glue robot at Manchester

MANCHESTER 1824

The University of Manchester

Assembly Continued

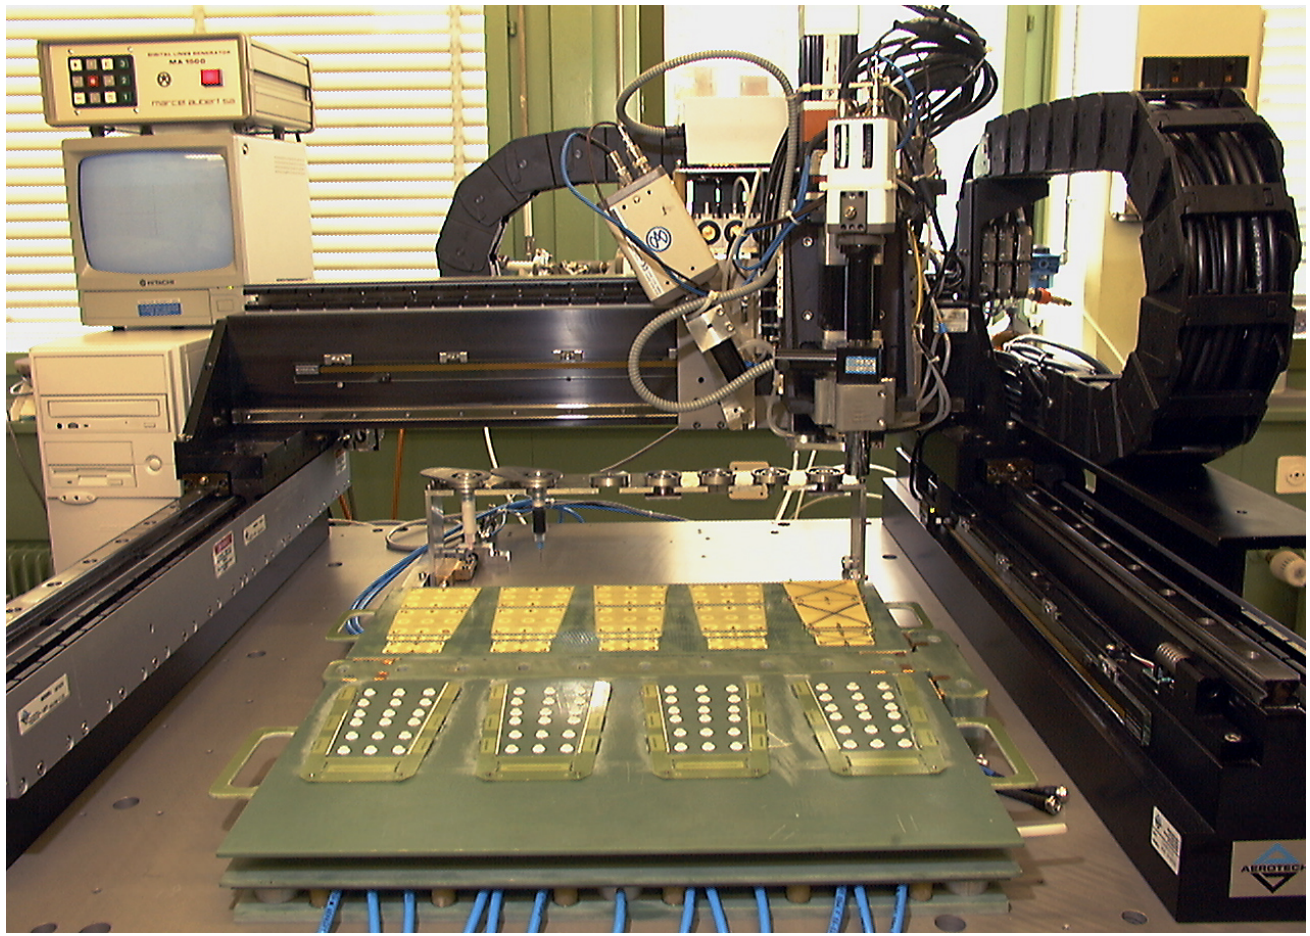

CMS tracker module assembly gantry

Assembly Continued

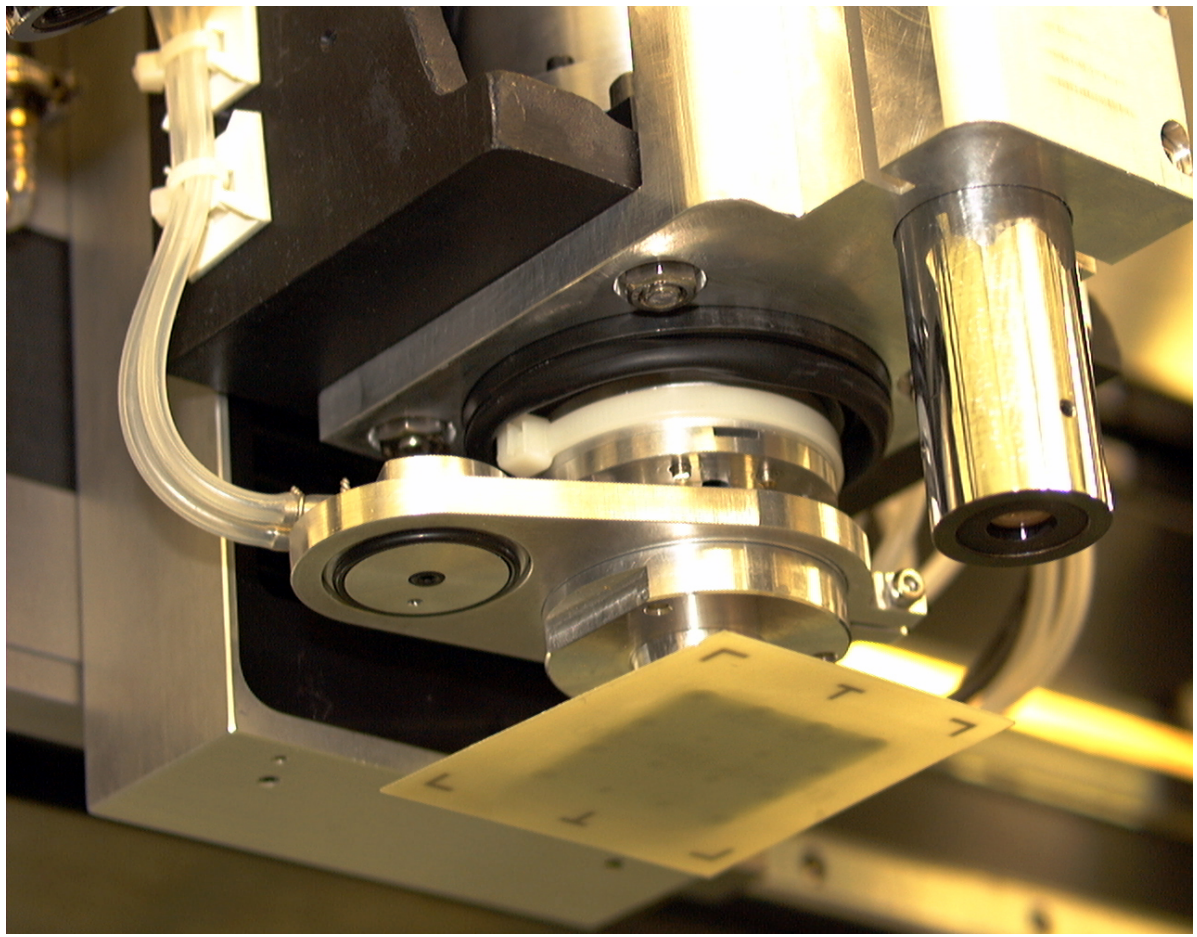

Interchangeable Vacuum pickup tool

High-Speed Networking

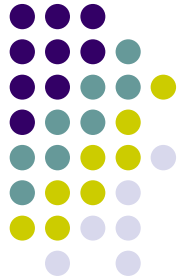

- Current work has focussed on PCI-X motherboards and cards
  - Measurements
    - Throughput
    - Packet-loss
- Plan to extend these studies to PCI-express
  - Repeat Ethernet throughput studies
  - Investigate use of FPGA (Virtex 4) as data source/sink at 10 Gig (both crate – ATCA – and PC based)
  - Look at redundant UDP routing using commercial switches (for routing data from “receivers” to processing farms)

10 Gigabit Ethernet: iperf TCP Intel Results

- X5DPE-G2 Supermicro PCs B2B
- Dual 2.2 GHz Xeon CPU FSB 533 MHz
- XFrame II NIC
- PCI-X mmrbc 512 bytes
- 1500 byte MTU
- 2.5 Mbyte TCP buffer size
- Iperf rate throughput of 2.33 Gbit/s

- PCI-X mmrbc 512 bytes
- 9000 byte MTU
- Iperf rate of 3.92 Gbit/s

- PCI-X mmrbc 4096 bytes
- 9000 byte MTU
- Iperf rate of 3.94 Gbit/s

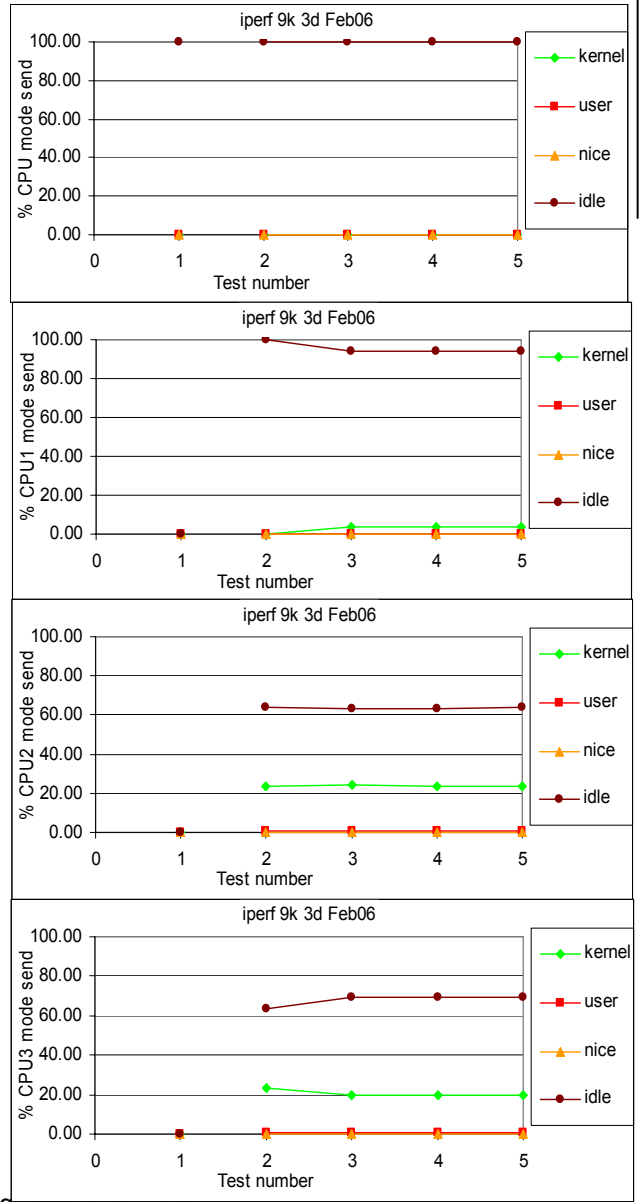

10 Gigabit Ethernet: UDP Intel Results

- X5DPE-G2 Supermicro PCs B2B
- Dual 2.2 GHz Xeon CPU FSB 533 MHz
- XFrame II NIC
- PCI-X mmrbc 4096 bytes

- Low rates
- Large packet loss
- ???

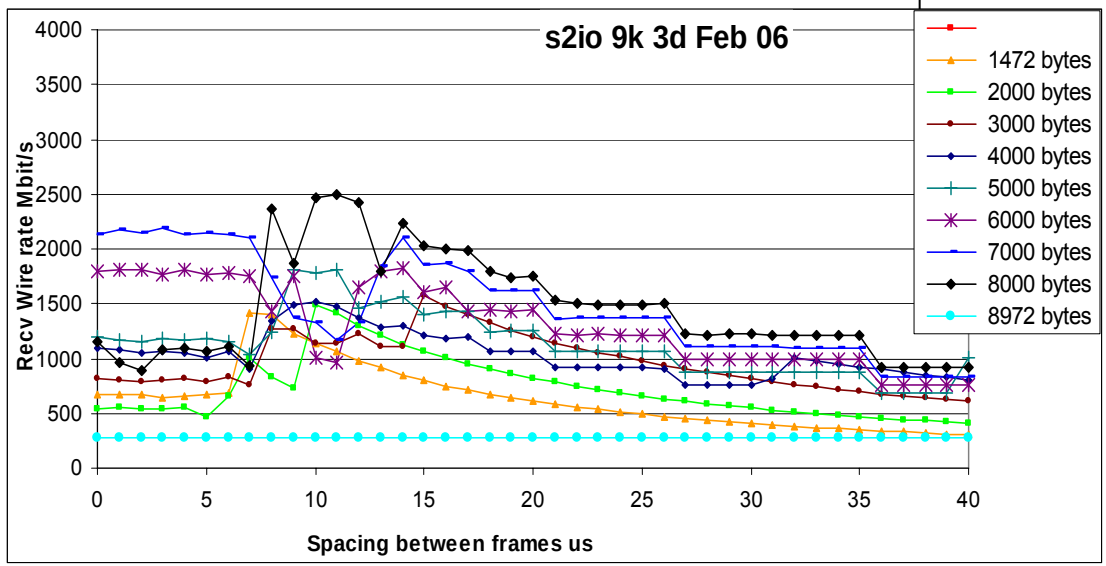

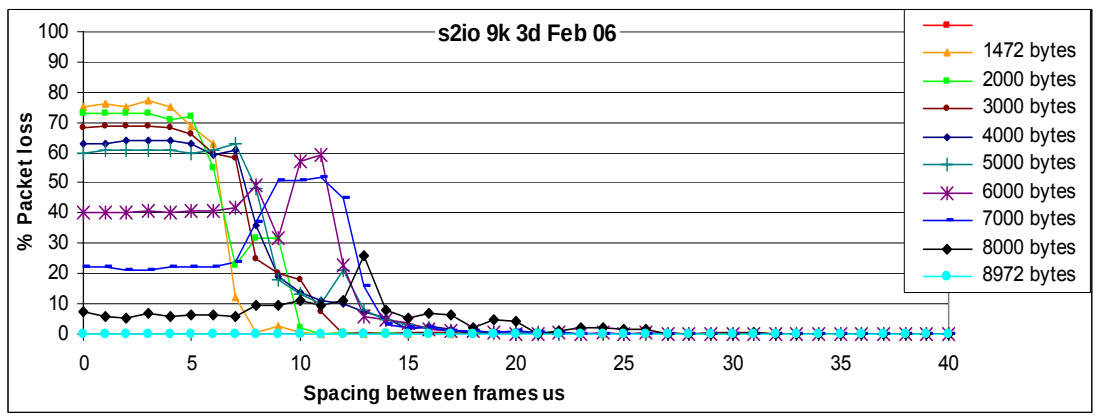

10 Gigabit Ethernet: TCP Data transfer on PCI-X

- Sun V20z 1.8GHz to 2.6 GHz Dual Opterons
- Connect via 6509
- XFrame II NIC
- PCI-X mmrbc 4096 bytes 66 MHz

- Two 9000 byte packets b2b
- Ave Rate 2.87 Gbit/s

- Burst of packets length 646.8 us
- Gap between bursts 343 us
- 2 Interrupts / burst

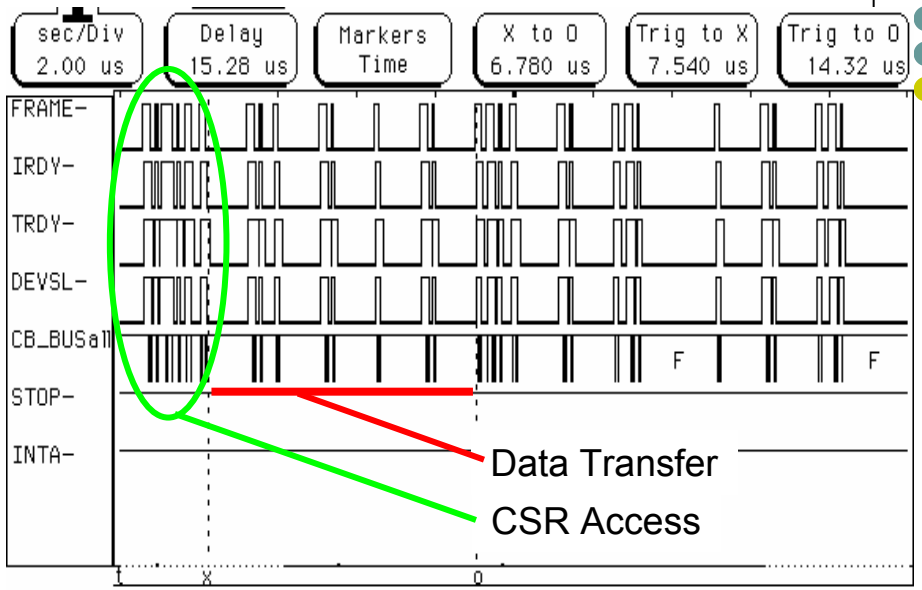

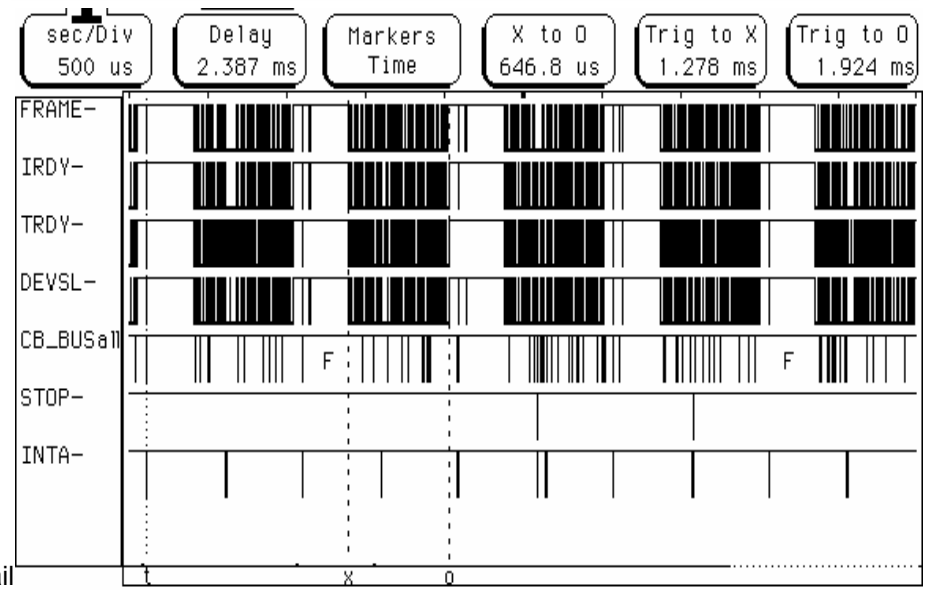

10 Gigabit Ethernet: UDP Data transfer on PCI-X

- Sun V20z 1.8GHz to 2.6 GHz Dual Opterons
- Connect via 6509
- XFrame II NIC
- PCI-X mmrbc 2048 bytes 66 MHz
- One 8000 byte packets
  - 2.8us for CSRs
  - 24.2 us data transfer effective rate 2.6 Gbit/s
- 2000 byte packet wait 0us
  - ~200ms pauses
- 8000 byte packet wait 0us
  - ~15ms between data blocks

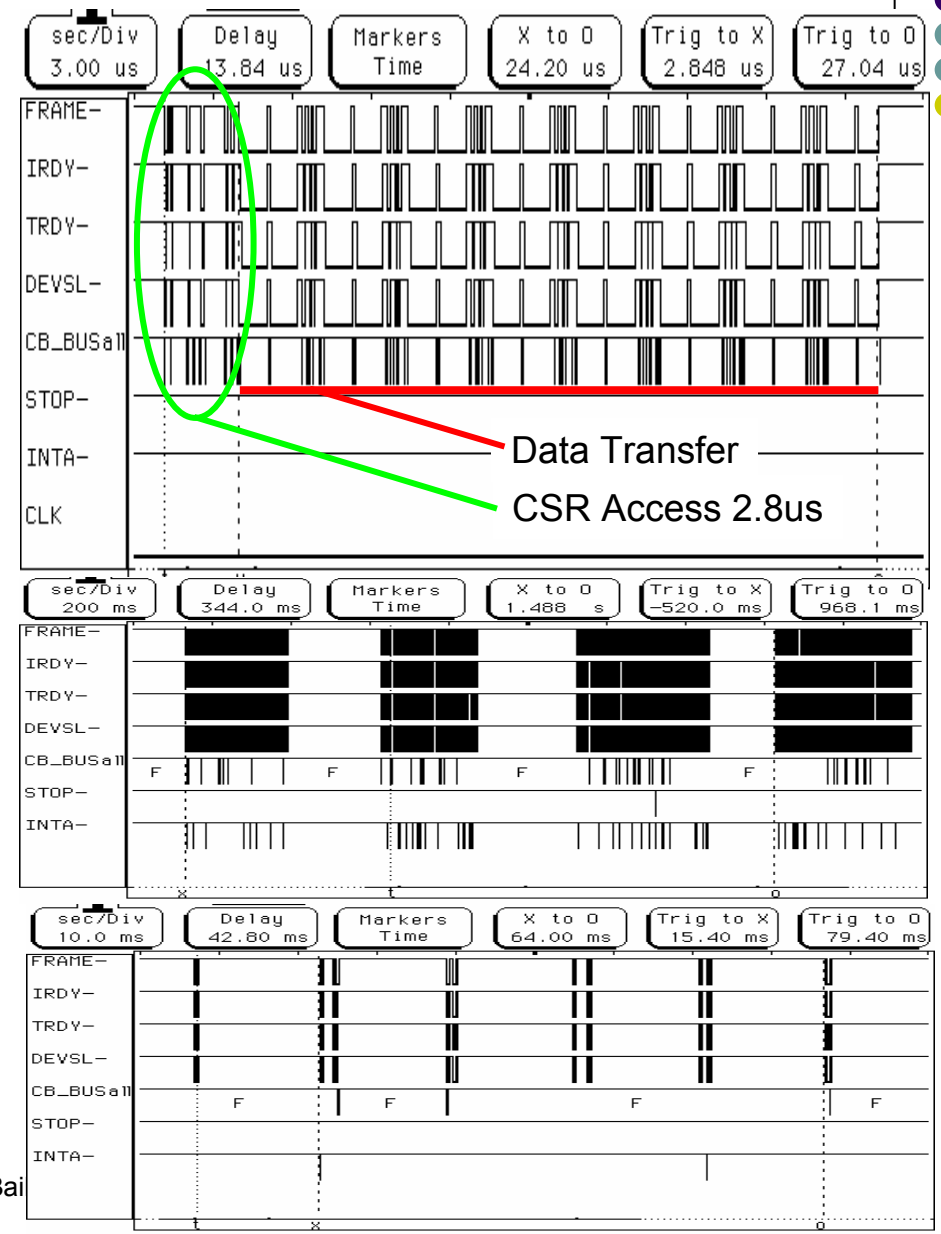

MANCHESTER 1824

The University of Manchester

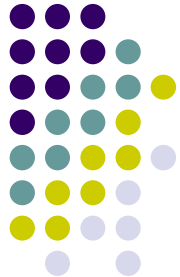

Conclusions

- Thermal simulations ongoing
  - Will build a cooling test rig in Manchester to verify simulations and act as a test bed for active cooling tests
- Mechanical
  - Glue tests now starting
    - Would like to investigate with “real” silicon
  - Have ideas for automated assembly based on ATLAS/CMS experience
- High-speed networking studies now being extended to PCI-express