2PAD Digital System

Architecture

Chris Shenton – JBCA - Nov 2009

MCCT-SKADS Technical Workshop Nov 2009

Strategic Objectives

- Demonstration of a functional fully digital beamformer
  - Small scale useable instrument – “2-PAD”

- Develop metrics to describe the primary characteristics of various technology platforms on which such beamformers can be built.
  - What are these characteristics?
    - Processing Capacity
    - Power Consumption
    - Cost
    - Scalability
  - As a function of (Bandwidth, Sensitivity, No of Beams, Beam Quality)
  - Feed upward to SKA system design team to allow objective design decision making

2PAD Design Aims (Digital System)

- Provide a modular platform for exploration of digital systems architectures for SKA

- Major Subsystems
  - Multi Channel Digitiser
    - *128 Channels @ 1GHz Sampling (500MHz Bandwidth)*
  - Clock Distribution Subsystem
    - *LVPECL 2 Level Clock Tree*
  - Channeliser
  - Beamformer
    - *Real Time Hardware*
    - *Quasi Real Time Software*
  - Back-end data processing
    - *High(ish) end Multi Core Server*
  - Management Subsystem
  - Infrastructure
    - *Cooling*
    - *RFI Safe Enclosure*
    - *Power Supply*

Technology Options

Candidate Technologies

- HPC Clusters
  - Commodity Hardware, easy access to technology, Upgradeable, Standards based Networking, Mature(ish) programming environment.
  - IO Bandwidth limitations, inter-processor bandwidth bottlenecks, power requirements, space & infrastructure requirements.
- Multi-Core processors (inc DSP/GPU/NP)
  - Semi-commodity hardware, maturing software tools, Accessible technology.
  - Immature technology, 'difficult' programming environment, space & infrastructure requirements, limited IO interconnect options.
- FPGA based reconfigurable computers
  - Can be optimised, medium power, medium cost, highly flexible. Good IO options.
  - Medium NRE Costs, medium power, medium cost. Specialist programming for high quality results.
- ASIC based dedicated processors
  - Optimised, lowest power, lowest cost (if volume is big enough)
  - Design effort, High NRE, Poor Flexibility.

MANCHESTER 1824

Technology Summary

	Infrastructure Requirements	Power	Flexibility	NRE Cost	Unit Cost	Performance (per Watt)
HPC	High	High	High	Medium	High	Low
Specialised Processors (MC,GPU,ALU Array)	Medium	High-Medium	High	Medium	Medium	Medium
FPGA	Low	Medium	High	Medium	Medium	Medium-High
ASIC	Low	Low	Low	High	Low	High

Architectural Features

What are the features which drive the system characteristics?

- **Processing Elements**
  - ALU's
  - Dedicated Hardware Resources (Multipliers, Adders etc)
  - Microcoded Engines or Unrolled Pipelined Datapaths
- **Inter-element Communications**
  - Shared Memory
  - Point to Point Connections
  - Datapath Duplication & Fanout
  - Switched Interconnects (Fabrics)
- **IO**
  - Off Chip communications

Trade-Offs

- We need to understand the effect of various optimisations in terms of quality of results.
- Can we make significant simplifications to the processing to reduce power and area? Without compromising the performance of the instrument.
- Process technology – do we really need to be at the bleeding edge? Is Moore's Law always the answer?
- What does flexibility cost and how much can we afford?
- T_{sys} trade-offs as a function of hardware simplification.
- Data Quantisation, scaling, 2^n Co-efficient Quantisation.
- Can we sacrifice T_{sys} in order to meet power & cost requirements?
- Can we live with sub-optimal BERR performance across interfaces?
- Requirements for the clock distribution at various levels of the system hierarchy.
- Clock tree's are potentially very power hungry.

2-PAD Simplified System Overview

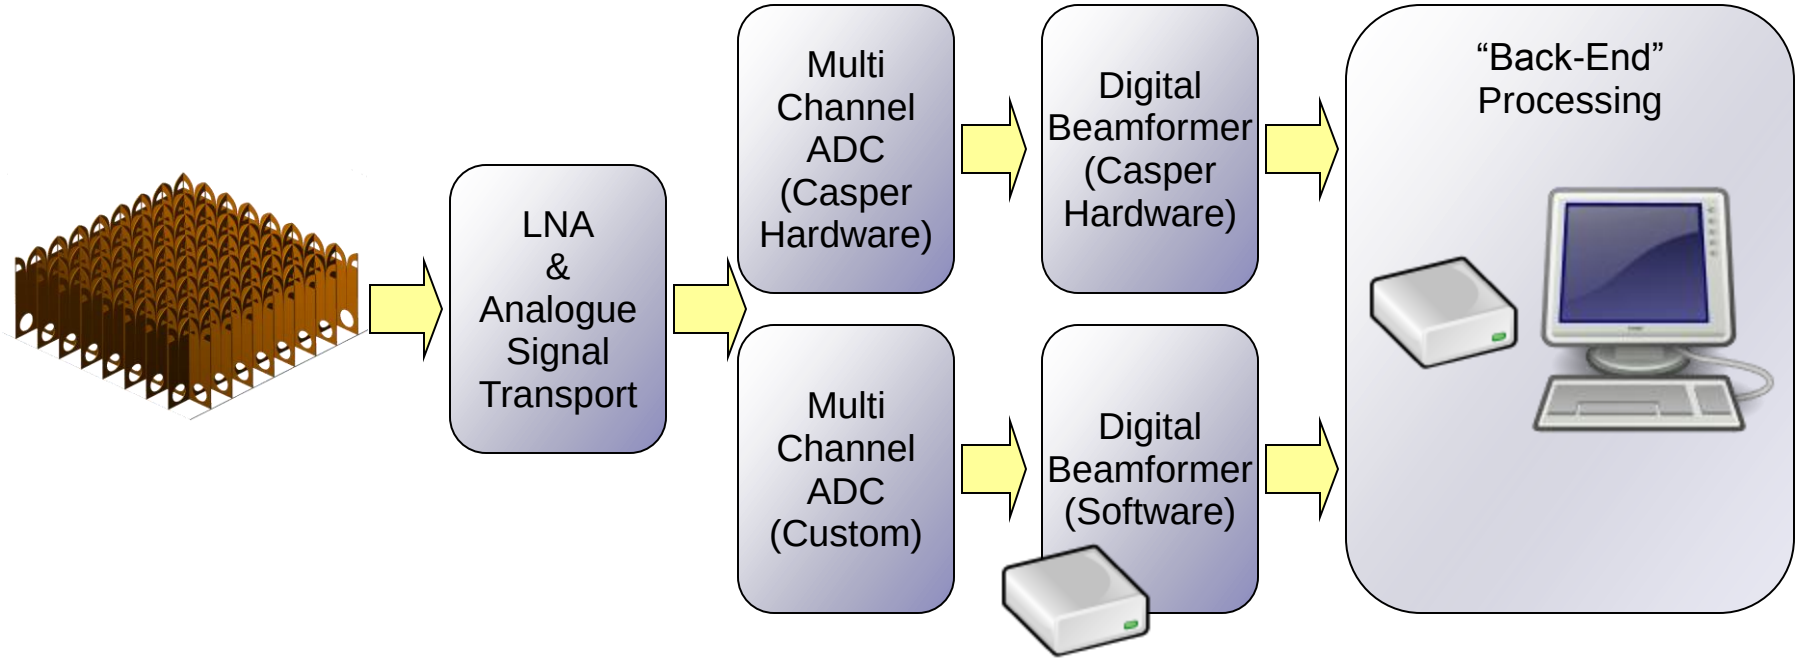

2-PAD Slightly Less Simplified Overview

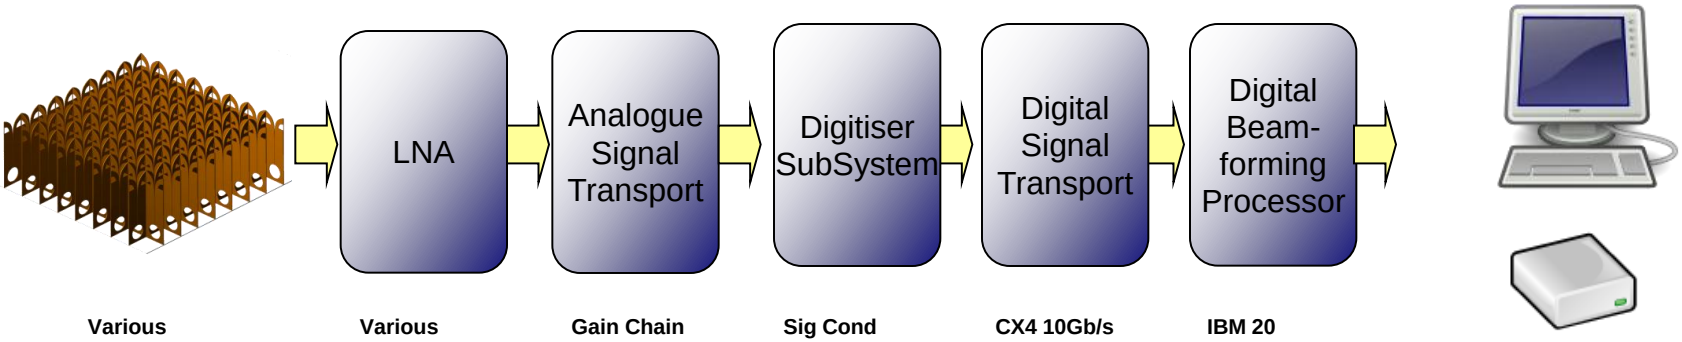

Various

BECA
ORA

Various

To suit Antenna Options

Gain Chain

CAT7 Cable

Sig Cond Card

Midplane

DAQ Card

Clock Distribution

CX4 10Gb/s

4x3.125Gb/s

8b10 Encoding

Streaming Channel

No Transport Overhead

IBM 20 Node C64 "Cyclops" Based System

DELL Quad Core ZEON Server

2TB Storage

RHEL

Digitizer

- 3 Main Elements;
  - DAQ Board – Dual Channel ADC + FPGA
  - Mid-Plane
  - Clock Distribution
  - Signal Conditioning Module

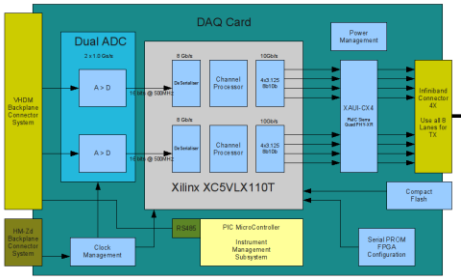

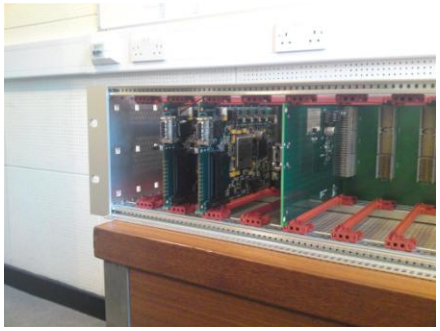

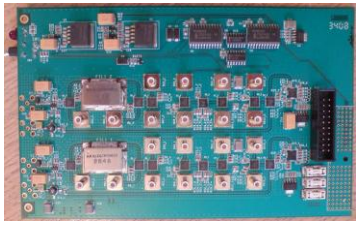

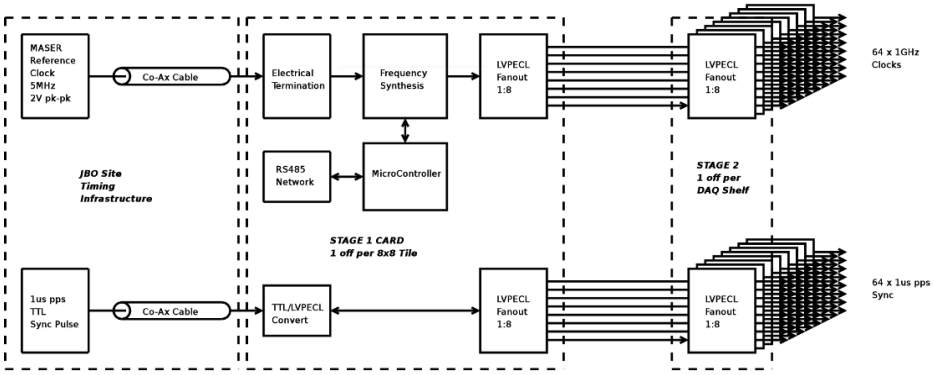

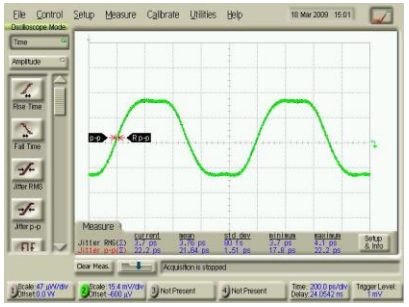

Physical

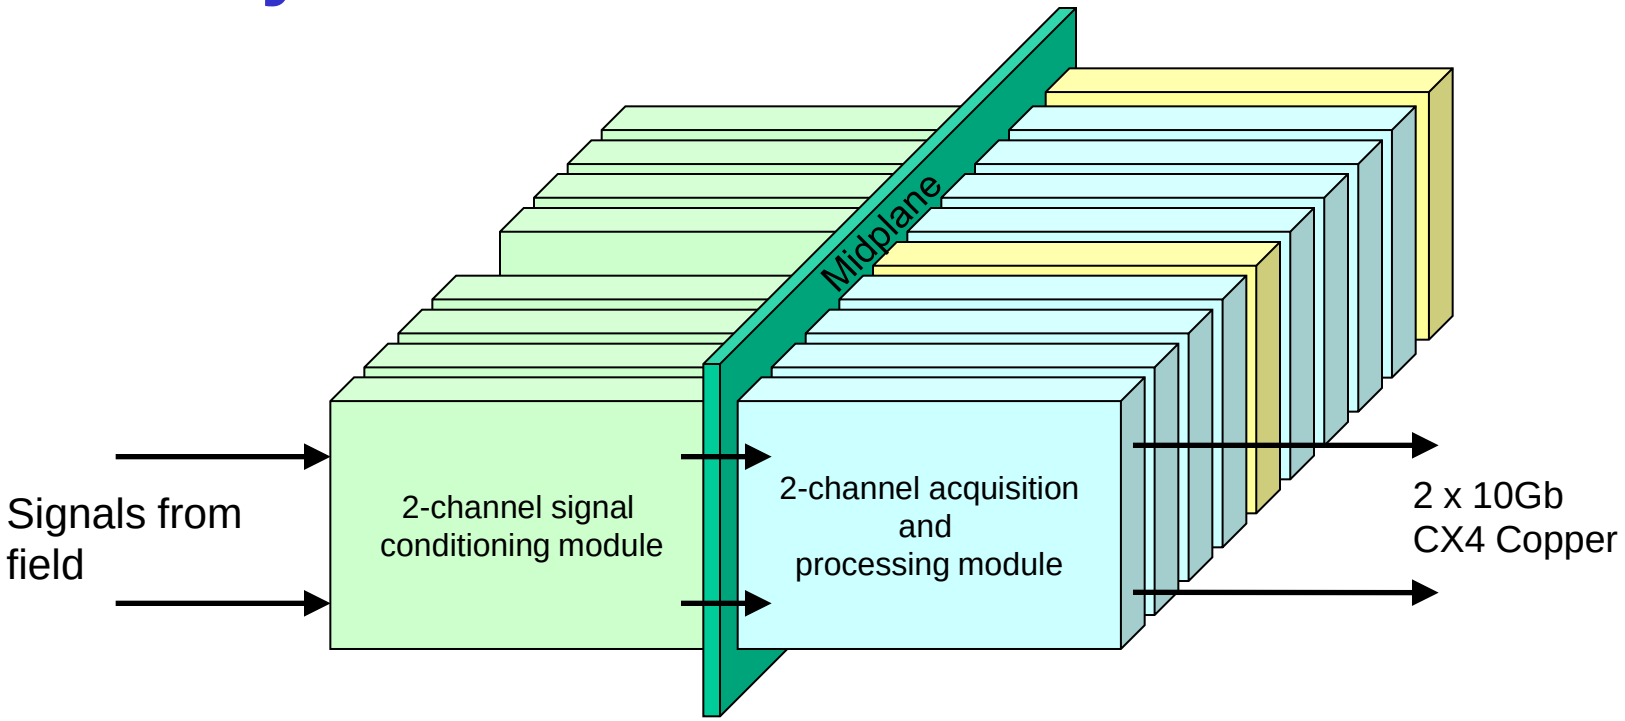

- Midplane based 3U 19” EuroCard Cardframe
  - Power & Clock Distribution via midplane
  - Instrument Management Network via midplane
  - Standard Prefab Commercial Metalwork
  - Nominally 10 Slots (8 DAQ, Shelf Clock, Optional System Clock)

ADC Performance

- Characterisation of the ADC and DAQ ‘front end’
- Tests carried out using RF signal generator including Signal Conditioning Module in signal path.
- Linearity and S/N figures for the DAQ/SCM combination under development… We need to understand how much we are degrading the analogue signal ‘quality’ through the digitiser.

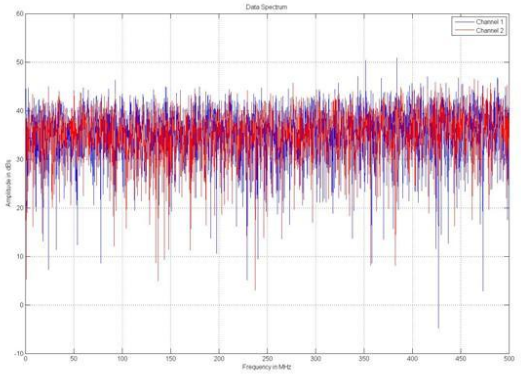

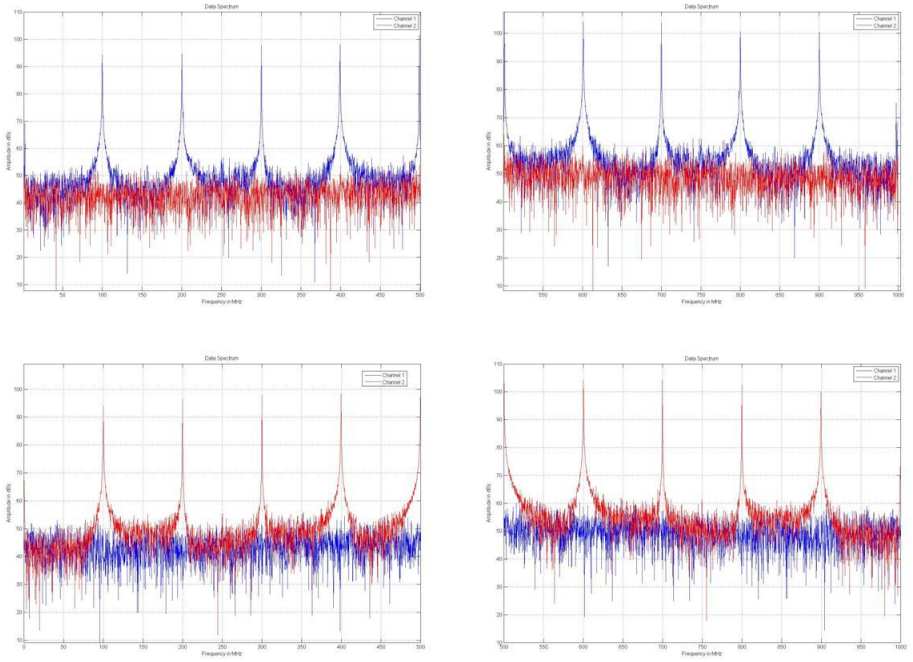

A Few Words About Clock Distribution

- Lots of it required for an array based system
- Very power hungry
- Hierarchical – But jitter is additive so what's the budget?
  - 1024 Sampled Channels needs 2^{10} Clock Endpoints
  - Typical LVPECL Clock Buffer 1:4
  - 5 Level Clock Hierarchy
  - 20 ps per level jitter
  - 100 ps Jitter
- Total Jitter **will** escalate rapidly with high 2^n clock trees

Progress in SKADS

- Digitiser subsystem – implemented and characterised
  - Work continuing in area of optimised frequency channelization
  - Bringup of more acquisition channels
  - Clock subsystem implemented and characterised
  - Fine grain control of acquisition timing & deskew
- Instrument Management Subsystem implemented and in use.
  - Supports TCL scripting
  - Speedup of interface planned
- Data transport using standards based approach
  - currently XGMII over CX4 (simplification from original Aurora design)
  - Support for 10GbE via iBOB… Work to implement native Tx Only 10GbE in progress.
- Software beamformer based on IBM Cyclops multi core processor
  - Hardware not at JBO but still available remotely at IBM location
  - Not Real Time but still meaningful
- Hardware beamformer & channelizer built using CASPER Hardware (See KZA talk)

2PAD into PrepSKA

- Modular approach allows “plug and play” of alternative implementations…
  … we can plug things together…
  … and play with it…

- Inherent capability provides ‘room’ for meaningful exploration of RTL structures (essential for ASIC migration)
  - Channelisation
  - Beamforming (Casper, Uniboard, COTS, Other)
  - Non-Imaging processors (Uniboard)
  - RFI mitigation techniques (Uniboard)
  - Power Analysis at RTL Level
  - RTL prototyping

- Continuation of multi core software techniques in collaboration with IBM based on knowledge acquired during SKADS
  - Can use 2PAD Digitiser subsystem to acquire high quality data sets for offline processing using software techniques
  - Can model quasi-real time behaviour using block data approach

- Interface into eMerlin
  - Validation of phase transfer activities in AA context.
  - Understanding of the wider system timing hierarchy

Conclusions

- There is not a single “RIGHT” answer.
- The movement of data (both intra-chip and inter-chip) is at least as significant as the raw number of processing elements.
- The retained “Quality” of the data through the processing hierarchy has a cost (e.g. how hierarchical can a system be before beam quality is compromised to an unacceptable level).
- What is the cost of flexibility? Where do we need it, and where can we achieve flexibility by other means, for example dynamic co-efficient insertion.
- We are moving away from opinion based arguments towards generating the evidence we need to inform the SKA design decisions and to enable us to design a system meeting the scientific requirements.
- 2-PAD and its derivatives will enable us to begin quantifying these factors.