Signal Transport And Networks

Chris Shenton
Jodrell Bank Observatory
University of Manchester

STaN CoDR June 2011

Custom or Proprietary Solution

What does that really mean?

"Well of course we all know that a custom solution is always the most expensive..."

- Actually a custom implementation built from commercial standard building blocks
- No new esoteric technology involved
- Aiming for the removal of unnecessary cost
- Full engagement with Industrial Partners

System Topology

- In SKA Phase 1 the majority of elements are within AA lo stations; 250 dish Rx's vs 560,000 AA RX Elements
- Majority of elements within 5Km of centre
- Majority of elements separated by only a few metres
- Maximum of 100Km reach required in Phase 1
- Dishes grouped in small 'clumps' with AA station close
- Phase 1 deployment 2016

System Characteristics

- Asymmetric data flows
- Tolerant of bit errors within data
- Synchronous data transfer required
- Tolerant of high latency
- Many 'end points' in the networks
- Fine Grained at element level (10G)

What Kind of Networks Terminate at an End Point

- Data transmission (Digital, Proprietary)
- Reference Phase transfer (Analogue, Proprietary)
- Synchronisation (Digital, Proprietary)
- Monitoring & Control (Digital, COTS, Ethernet)

What Are We Aiming For?

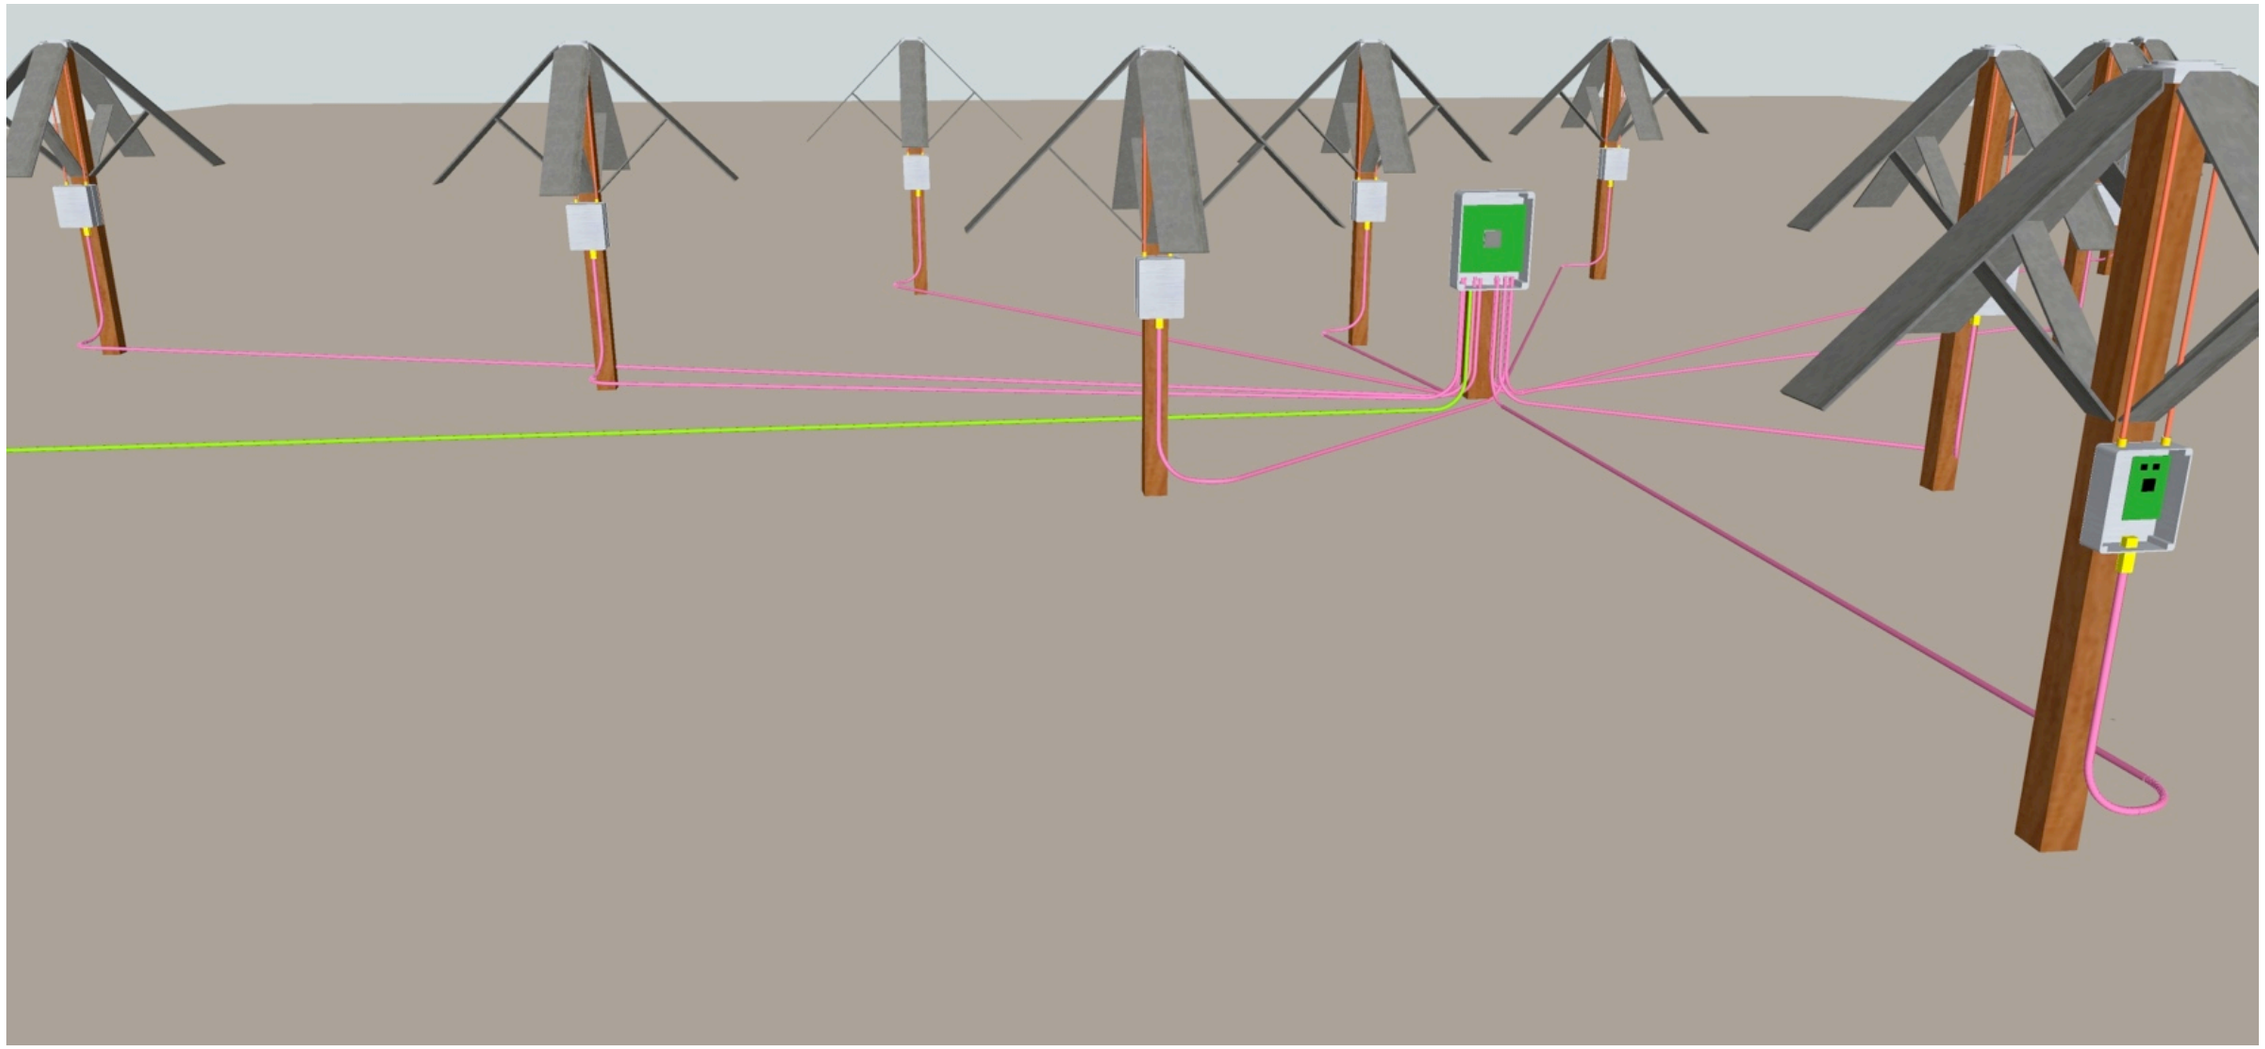

Something that looks a bit like this will dominate the system cost...

Physical Realisation Concept

- High Volume Manufacture > 0.5 Million units required...
- Design for low Cost
- LRU for simple maintenance and installation
- Reducing number of connectors improves reliability
- Not high tech
- Good enough... not as good as possible. It’s a Mini not a Rolls Royce
- Same design works for dishes

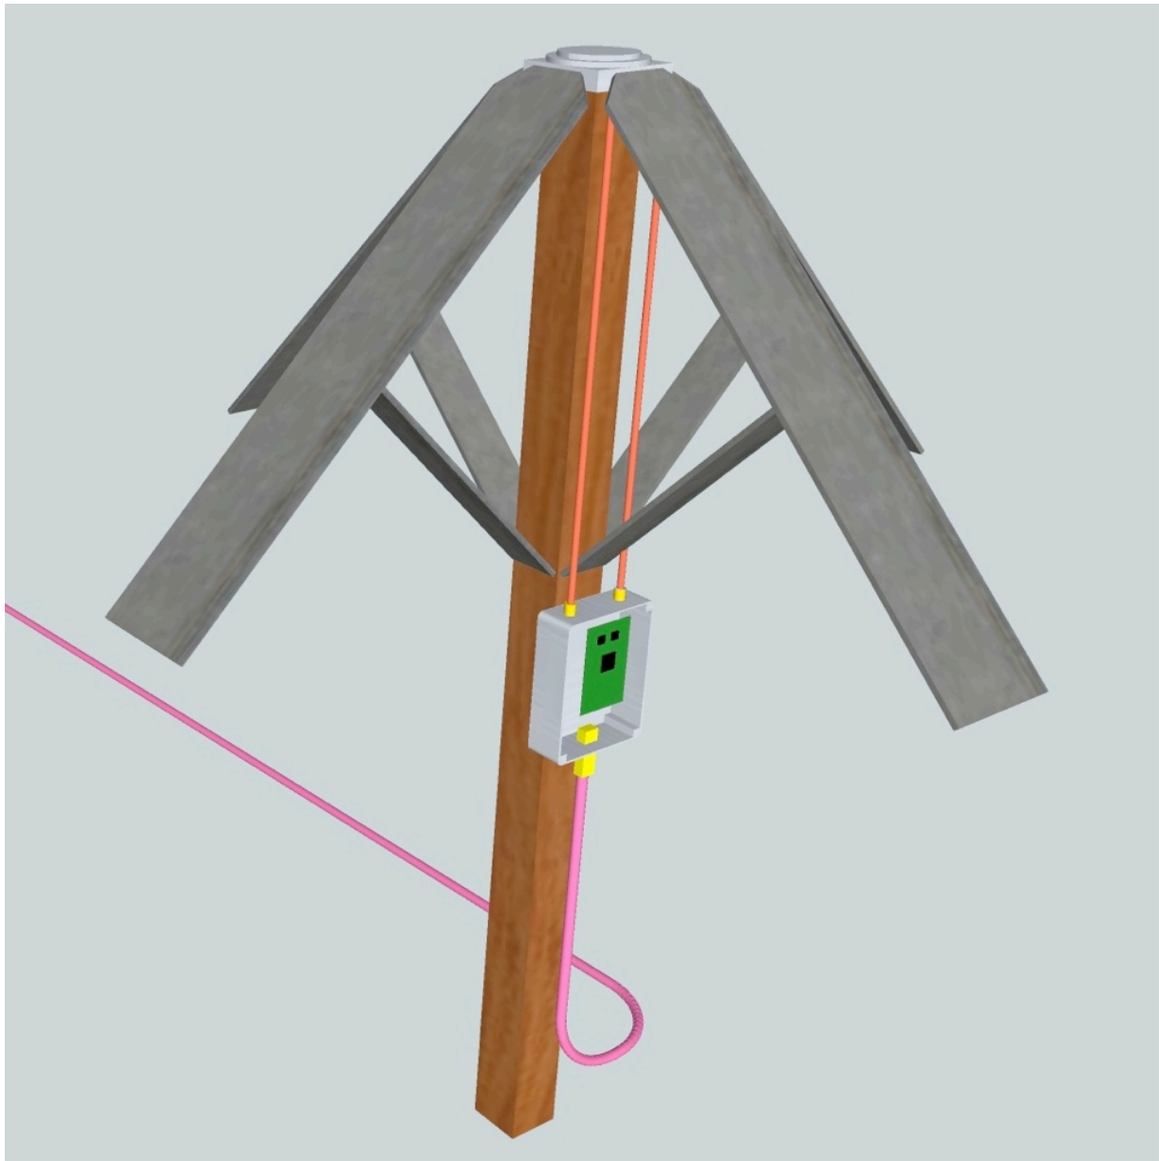

Integrated Digital Receiver Concept

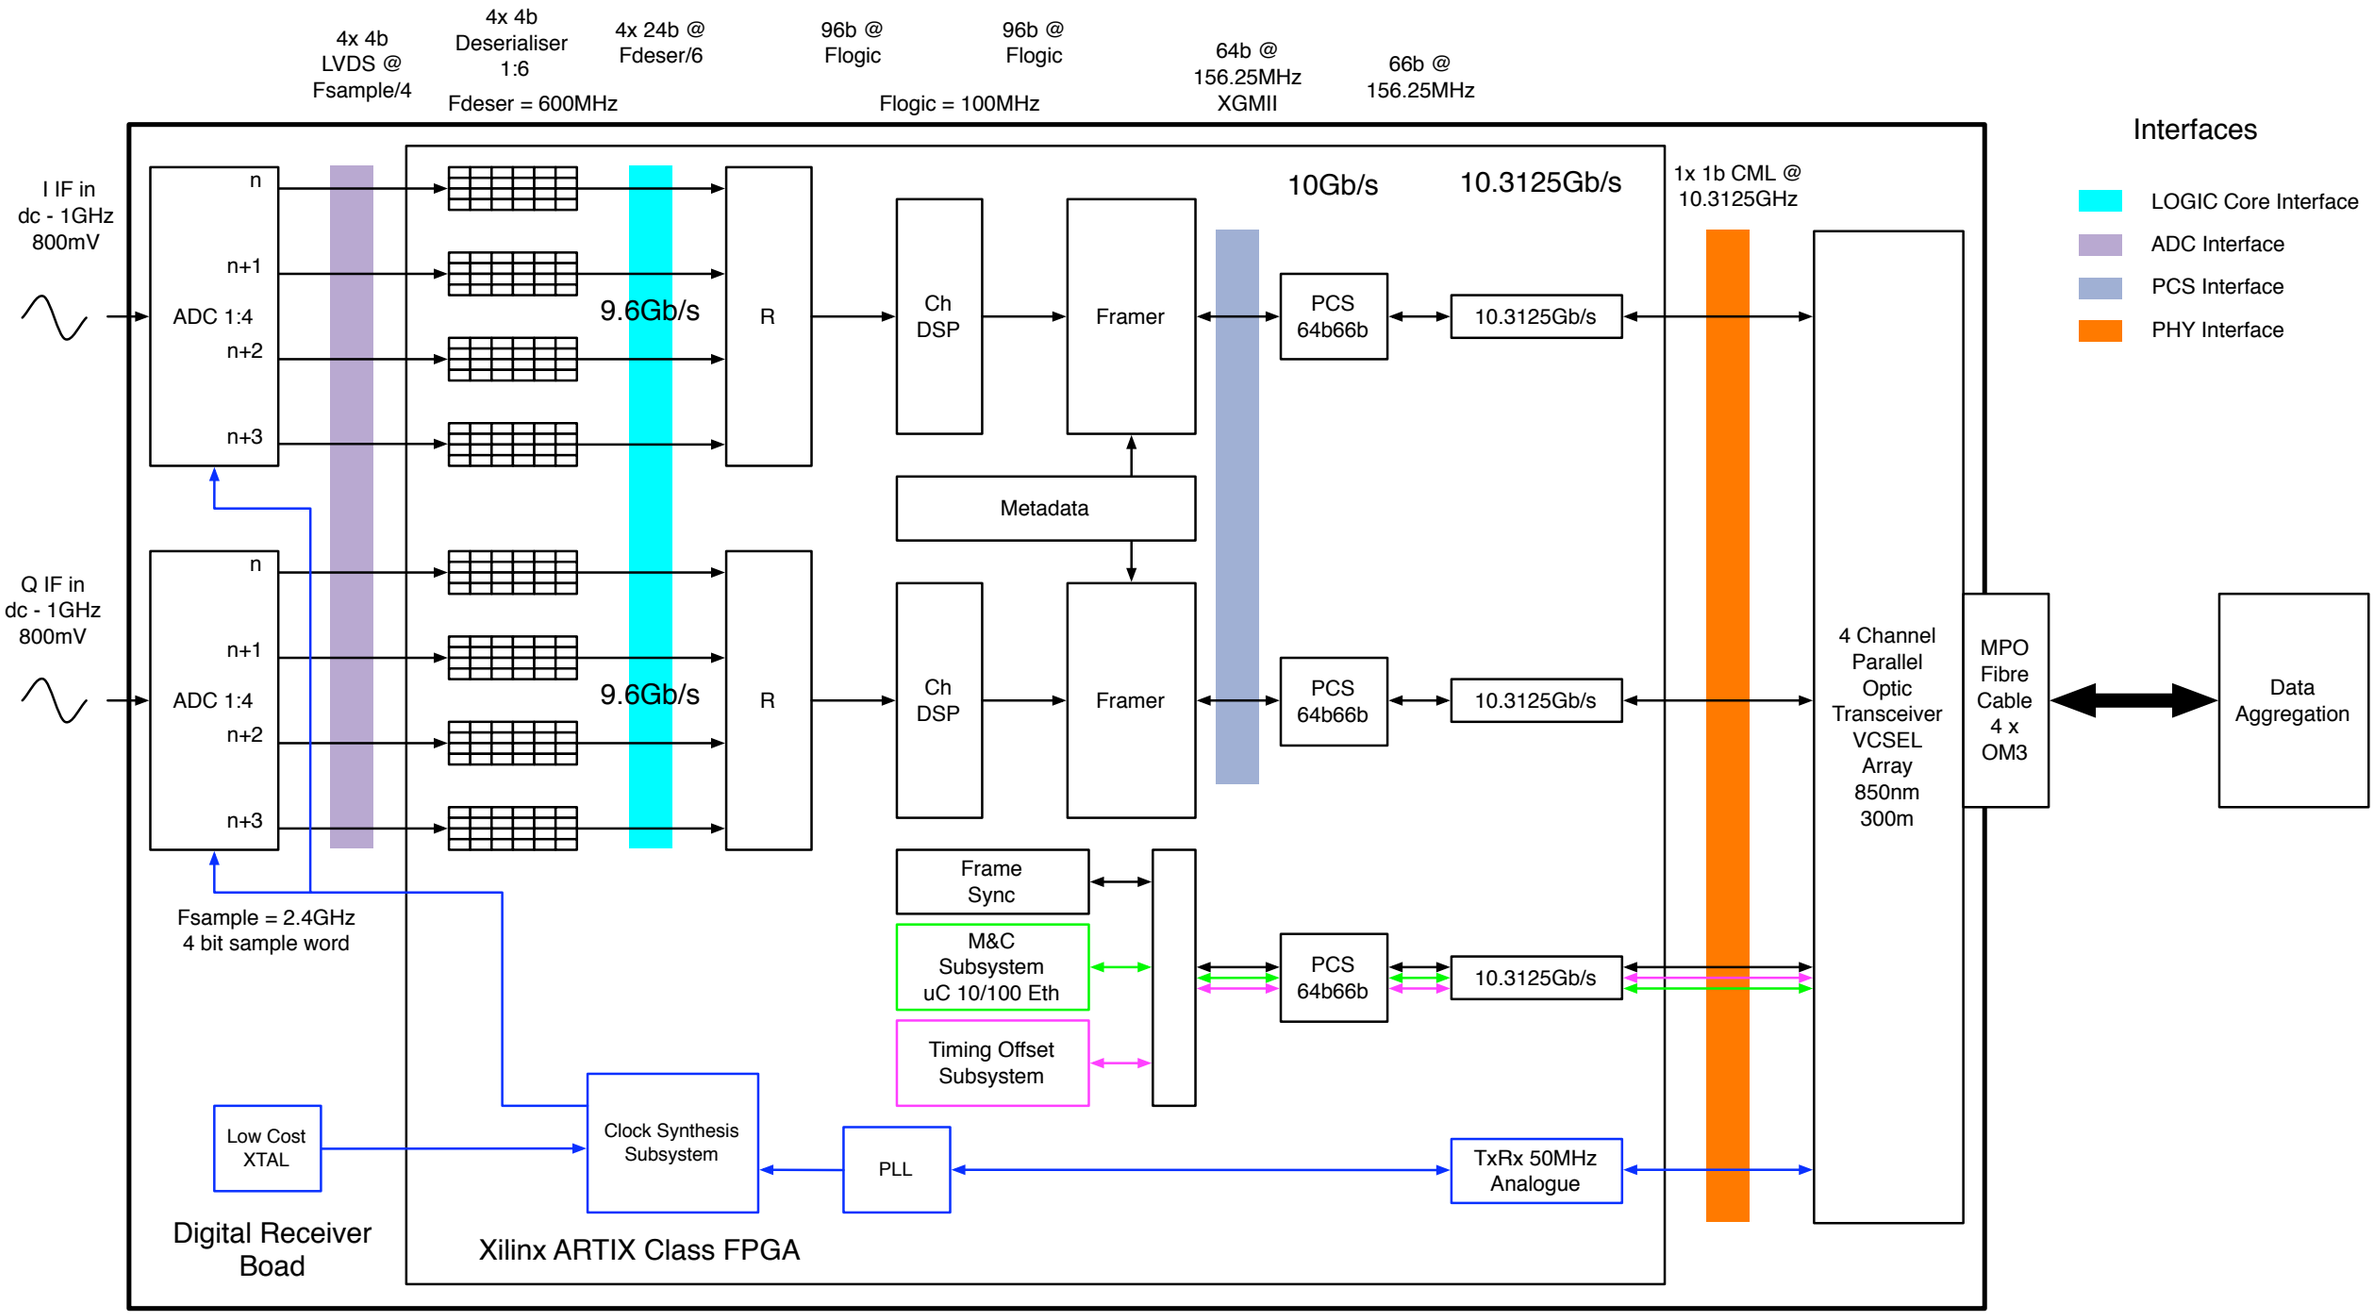

How to Reduce Costs

- Electrical & Logical Aggregation maximises the utilisation of available bandwidth
- Sharing of high cost Physical Resources such as interconnect and FPGA/ASIC platform
- Optical Aggregation via WDM maximises utilisation of physical fibre
- Standards based technology allows interception of commoditised technology
- Reduce the number of high cost terminations
- Reduce data volume by 'upstream' processing (this also has the happy side effect of simplifying the station and core processing)
- Reduce number of cables at Station and Core processing (1120 4x,600 12x)

System With Aggregation

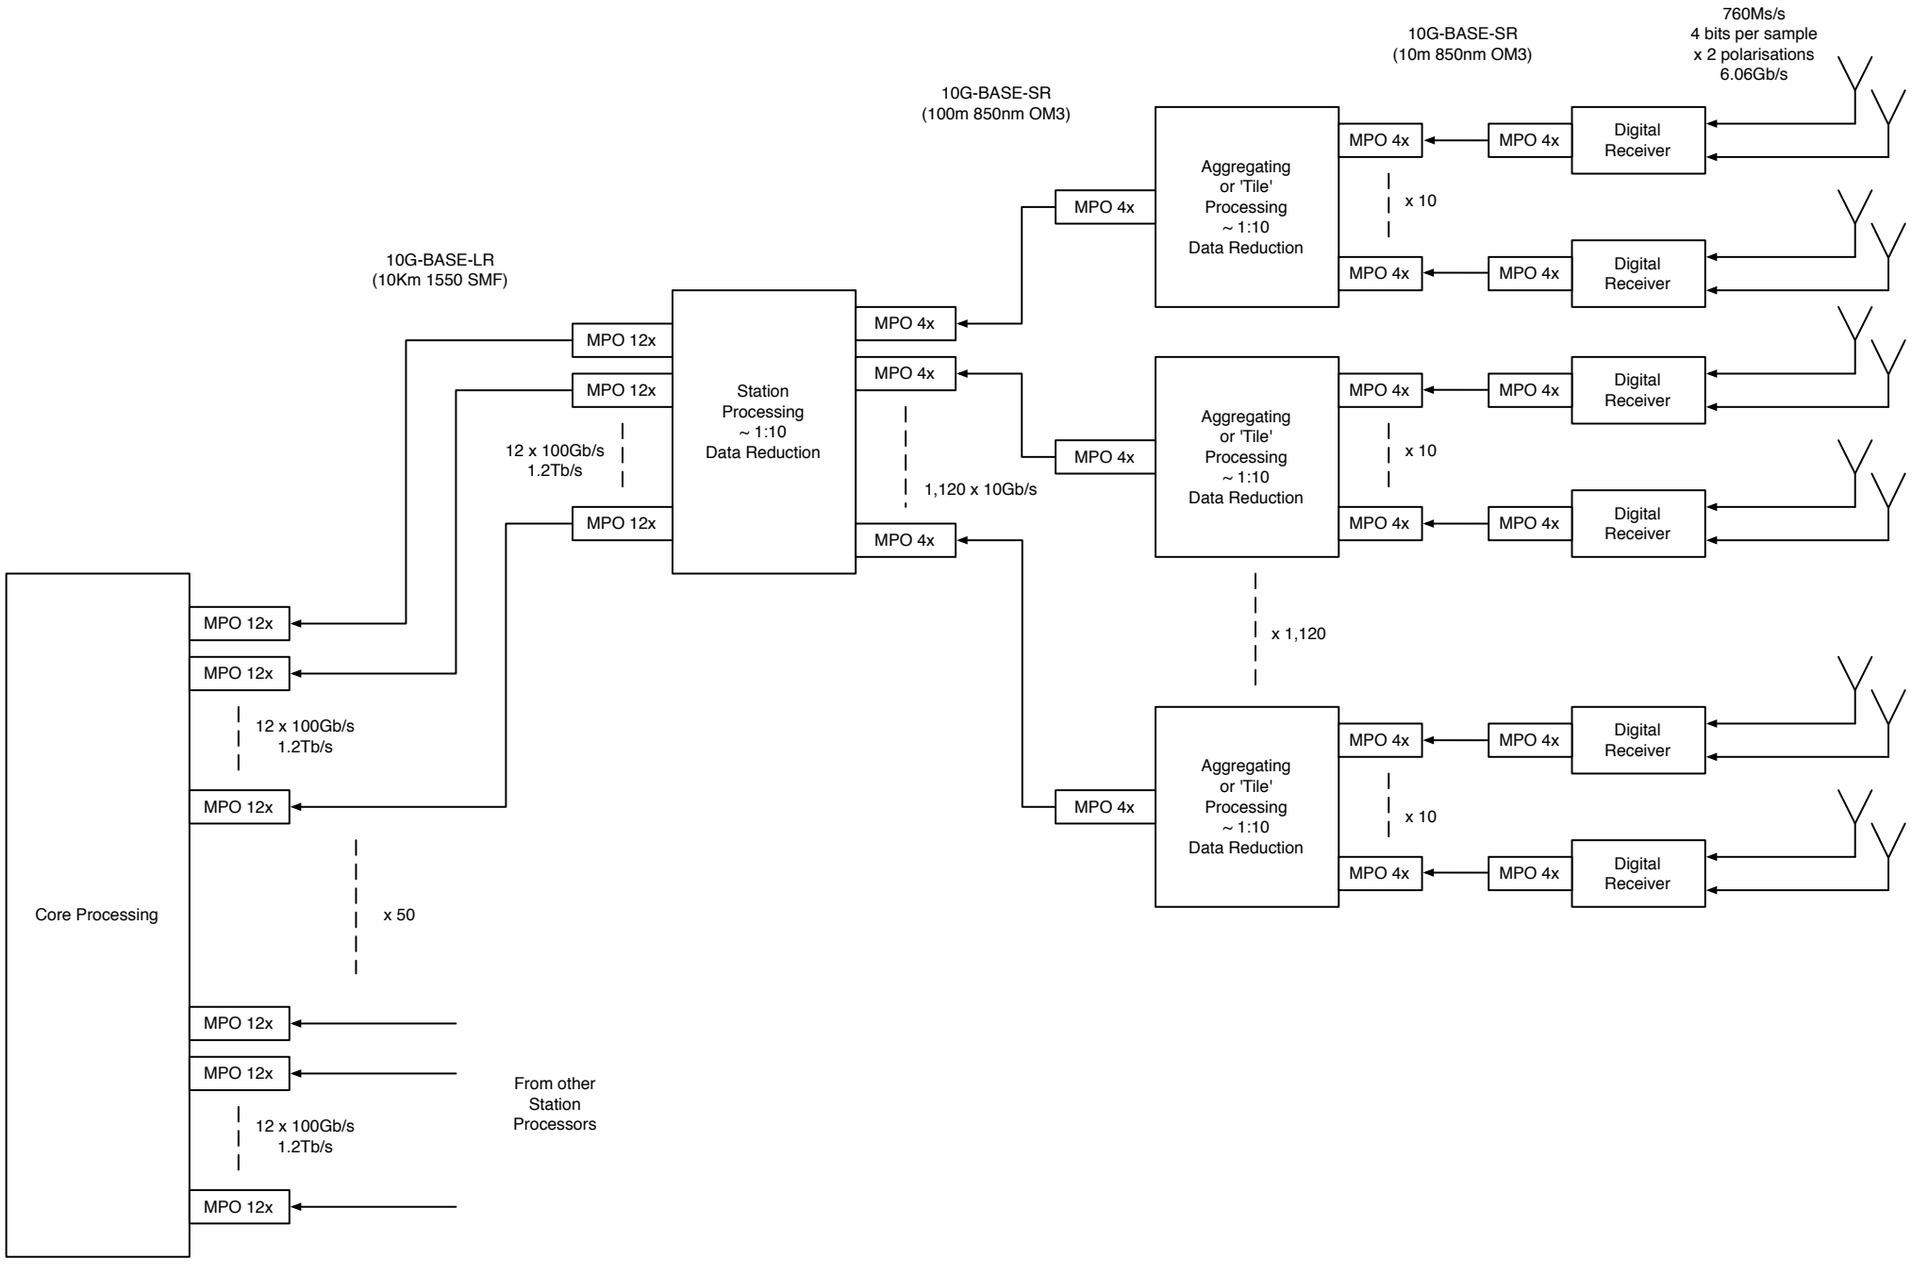

System Without Aggregation

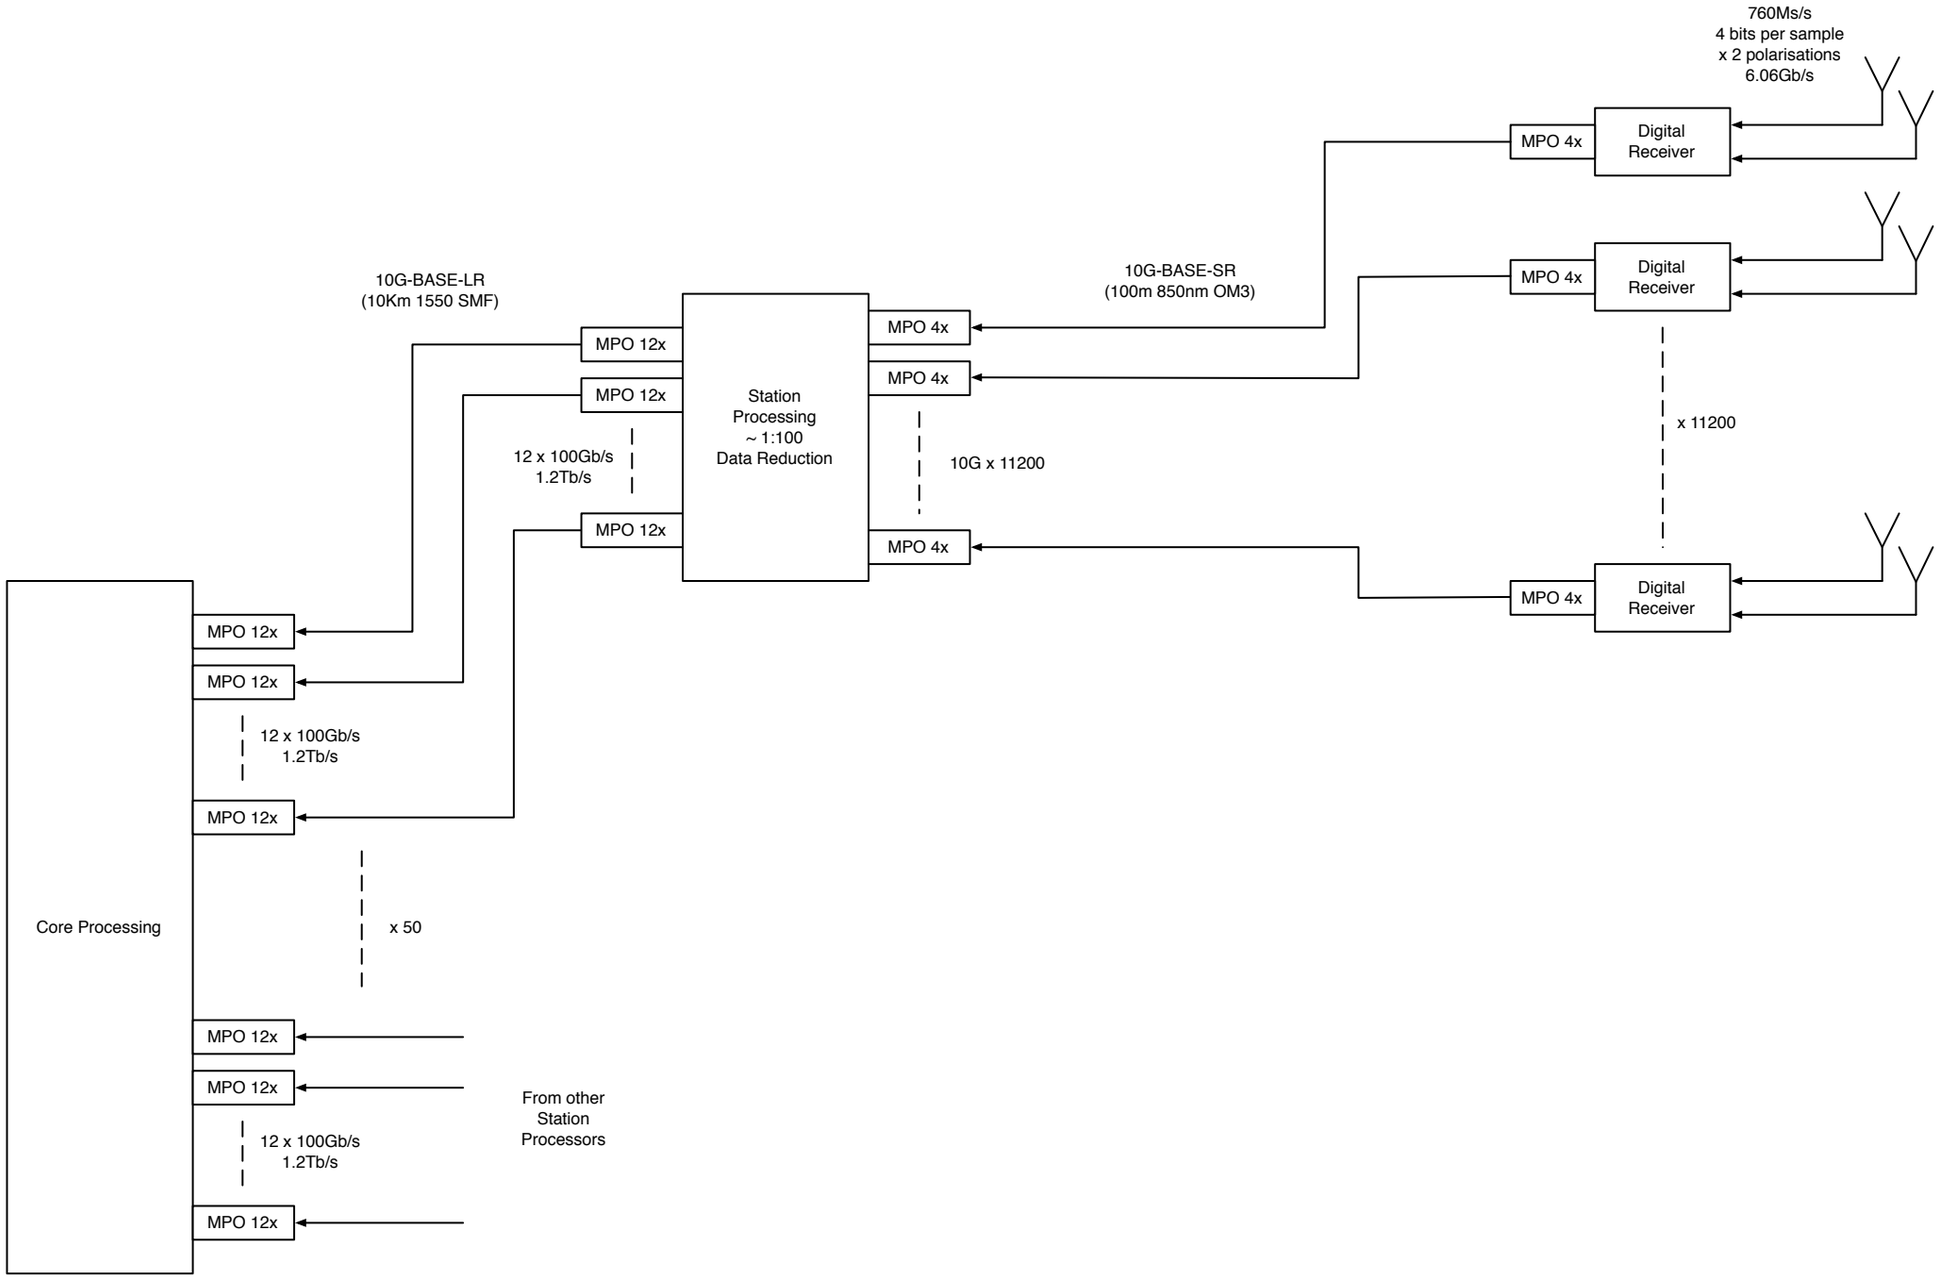

Strategy for getting to a Cost Effective System Solution

- Opportunities for resource sharing across multiple logical networks
- Use of appropriate technology for the application
- Obtain reduced costs by economies of scale
- Highest tech solution not necessarily appropriate
- Targeted use of high cost resources

Aggregation Processor Concept

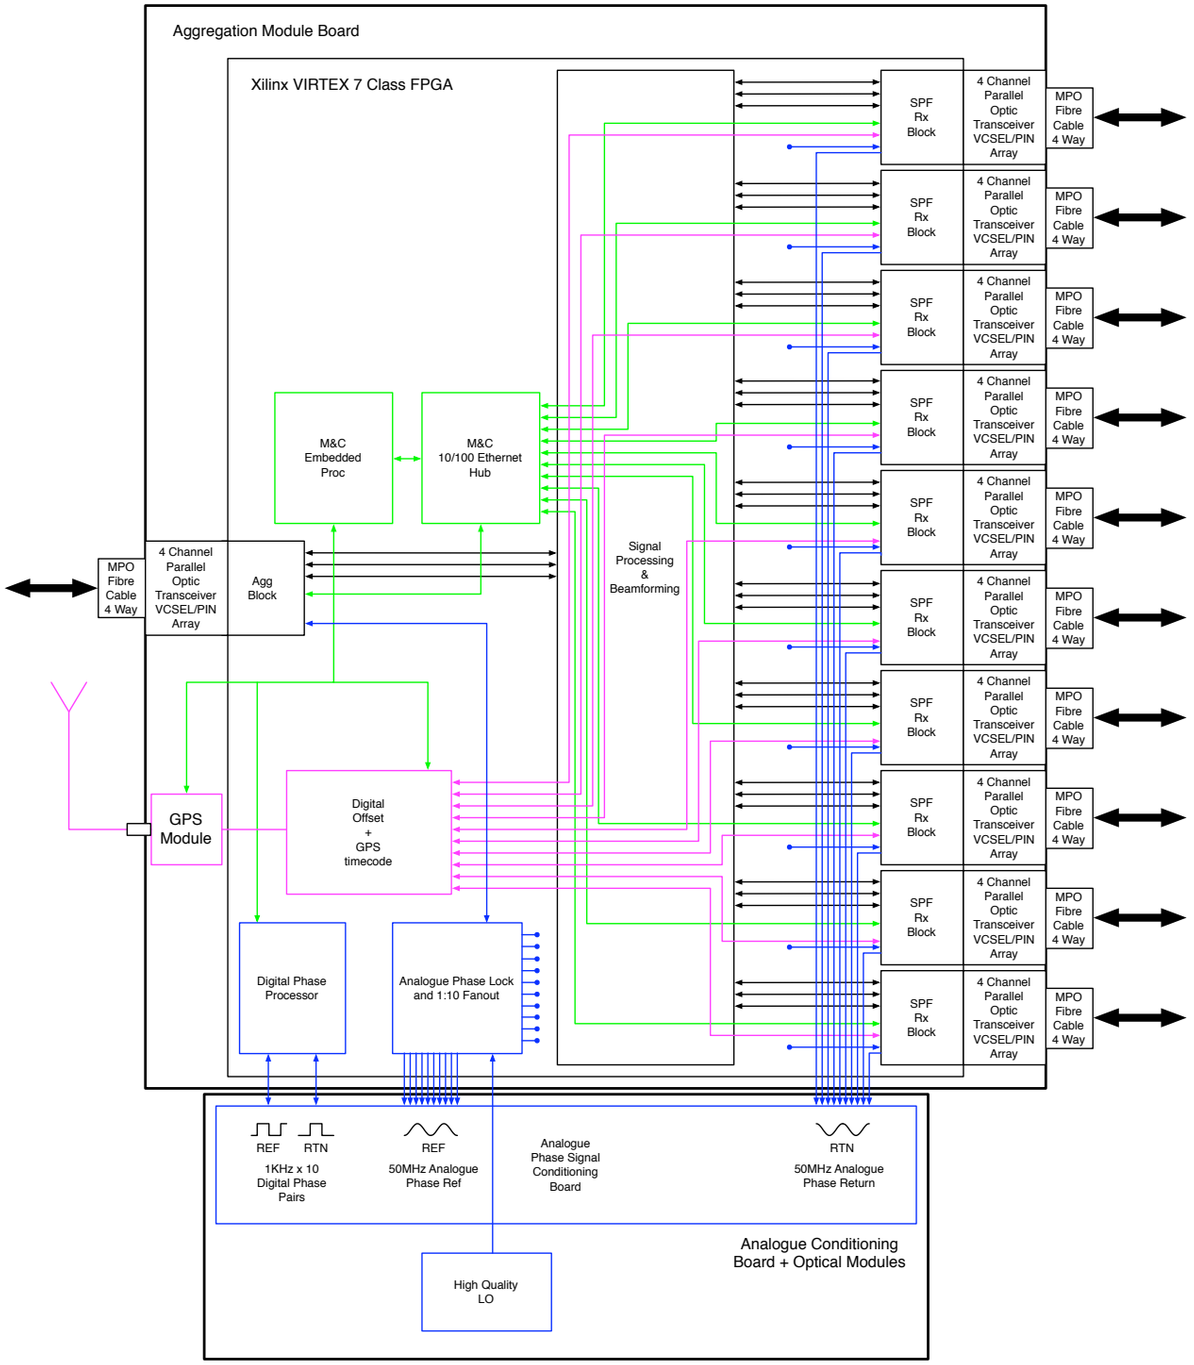

Electrical Aggregation and Tile Processing

- Majority of terminations short hop (560,000 connections less than 10m)
- Creation of sub groups or tiles allows electrical aggregation using low cost silicon ($10 class FPGA not $1000 class FPGA)
- Element Volumes make a simple ASIC a further cost reduction opportunity (cheap and cheerful 90nm technology will do this job... no need for esoteric 22nm...)
- Distribution of the processing into the end points and aggregators enables the maximum utilisation of logic resources.
- Major cost driver is IO in the form of the network termination
- Appropriate but simple processing at the element or tile level becomes effectively 'for free' as part of FPGA or ASIC logic
- Channelisation and forming of partial beams reduces upstream data rate

Physical Layer

- 4x/12x/24x 10G MPO Structured Cabling System
  - VCSEL Driver Array
  - 40G/120G/240G Capacity over fibre pairs
  - 25Gb/s VCSEL TXCR's in development pushing capacity to 100G/300G/600G
  - 850nm MMF (1310nm and 1550nm SMF products in development)
  - OM3 Cable for SR Links
  - IEEE 802.3 10G-BASE-SR/LR PHY specification
- Driven by data centre patch interconnect market
  - Target $0.25 per gigabit per termination ($20 per 4x)
  - Power Consumption 700mW per end 40% Tx/60% Rx per 4x (4 x 10G)
- MSA (SFP+) plug-able interfaces high cost and not getting cheaper quickly

TE Low Cost 10G-BASE-SR Solution

real & demonstrated - not slideware

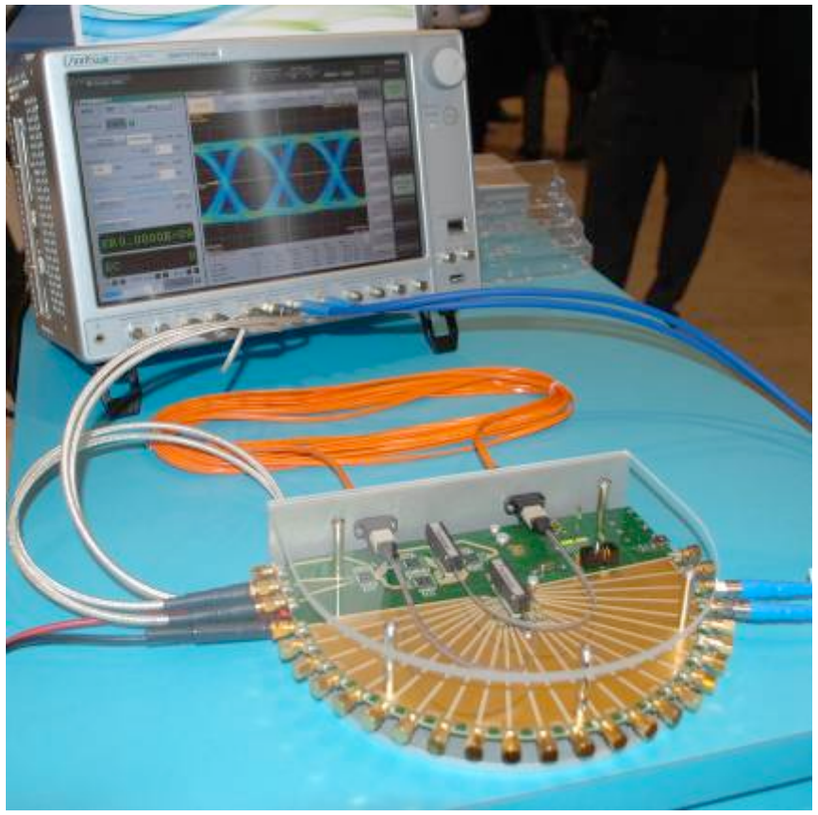

10G Now

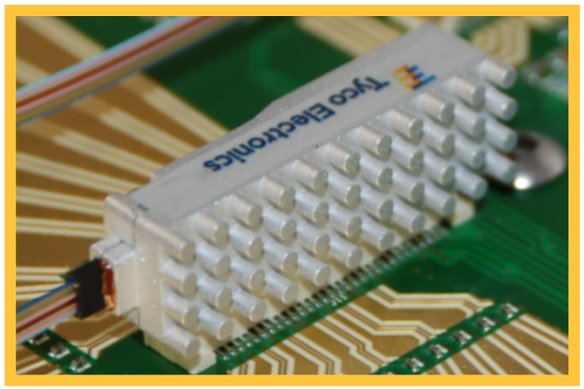

25G Target

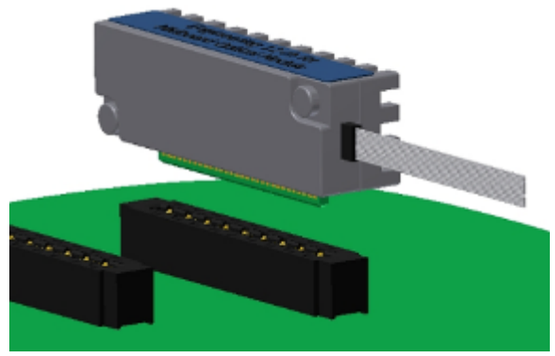

Transport Layer

- Proprietary
  - Simple Synchronous (Plesiochronous) framing required
  - Unidirectional
  - Minimal transport overhead required probably just a CRC
  - No retransmission, use it or lose it
  - Half as many driven Tx nodes in the data backhaul
- ‘Ethernet’ UDP/TCP/IP
  - Could be used for framing...
  - No real benefit
  - Costs Tx power
  - ‘Only really required for M&C network.
  - M&C Network over TCP/IP can easily be carried within synchronous frames as a VC
  - Could implement partial Layer 2-4 stack (ARP, Ping etc) at the endpoint... but why bother when it isn’t a switched network?

Can We Use 100Gb DP-QPSK Technology to Reduce Costs

- Depends on cost of transceiver technology - 10G SFP+ still expensive and cost reduction curve looks shallow
- Could reduce number of connections from Station to Core processing
- Still immature technology
- Probably required in SKA 2 for dense AA

Use of DWDM

- Optical Aggregation sub assemblies are expensive
  - Optical MUX's and DeMUX's at 50GHz channel spacing very expensive
  - Relationship between number of EDFA's and no of channels per fibre needs analysing for best cost tradeoff
- Point to point many core fibre is becoming more widespread
  - Driven by 'Fibre to the Home'
- Use of DWDM could be cost effective in long reach links with many stations
- Deploy where physical fibre costs become significant proportion of system cost
- Cost Tradeoff analysis needs to be done to find optimum solution
- Potentially more significant as the number of remote stations increases in Phase 2

Interface Standardisation

- ADC ➡ Logic
    - LVDS (or low power derivative)
    - Lower power solution required for dense arrays
- Electrical TXCVR ➡ Optical Module
    - XGMII (or CGMII if 100G arrives)
    - Plug-able Module?
        - SFP+ is still very expensive
        - SFP+ Cost includes high percentage of housing and mechanical assembly
        - Integrate the E/O and O/E into the Digital Receiver subassembly
            - Reduces cost
            - Improves Reliability

Cost Breakdown Aggregated

Based on 10:1 Aggregation

A 10 Element Group Consists of;

10 x DR FPGA/ASIC ($10)
10 x 10m Terminated Optical 4x MPO Cable Assembly ($30)
1 x AGG FPGA/ASIC ($100)
1 x 100m Terminated Optical 4x MPO Cable Assembly ($120)

Total Cost per Group

100+300+100+120 = $620 per group of 10 elements

A 11200 Element Station Consists of;

1120 x 10 Element Groups ($620)
12 x 100m Terminated Optical 4x MPO Cable Assembly ($120)

Total Cost per Station (at station inputs)

1120 x 620 = $ 694400 per station

Station to Core Connectivity Cost
12x MPO Termination $60
Fibre 5Km SMF in 192 cores @$5 per metre
24 Cores $3125 per 12x Connection
Fibre termination (Shared Fibre Cable) $800

Total is;

$60 + $3125 + $800 = ~$4000 per connection

600 connections at the core; (12 x 50)
600 x $4000 = $2.4M

Total SKA 1 Cost

50 x $700K + 600 x $4K = $37.4M

Cost Breakdown Direct Connection to Station

No Aggregation

A Single Element DR Consists of;

1 x DR FPGA ($10)
1 x 100m Terminated Optical 4x MPO Cable Assembly ($120)

Total Cost per Element

10 + 120 = $130 per element

A 11200 Element Station Consists of;

11200 x Single Elements ($130)

Total Cost per Station (at station inputs)

11200 x 130 = $1.456M per station

Station to Core Connectivity Cost
12x MPO Termination $60
Fibre 5Km SMF in 192 cores @$5 per metre
24 Cores $3125 per 12x Connection
Fibre termination (Shared Fibre Cable) $800

Total is;

$60 + $3125 + $800 = ~$4000 per connection

600 connections at the core; (12 x 50)
600 x $4000 = $2.4M

Total SKA 1 Cost

50 x $1.456M + 600 x $4K = $75.2M

Power Requirements

Aggregated

11200 x (1.4W + 5W) = 72KW
1120 x (1.4W +20W) = 31KW

Total 104KW per station

Non Aggregated

11200 x (1.4W + 5W) = 72KW

Total 72KW per station

Risks

- Compromises are required in order to simplify and cost reduce the system
  - Hard to pin down spec from scientists, let alone impose constraints
    - Cost constraint now hard to ignore
- Data Rate estimates are questionable for AA technology
  - LOFAR has data growth in beamforming to maintain S/N performance
    - is the required station data rate actually bigger than current spec states...
- Proposal based on emerging technologies - pricing is fluid
  - We really need a cheap 10G-BASE-LR (10Km) option, 1550nm VCSEL based technology not currently mature
  - Work with technology providers, use their expertise to get real world solutions
- Subject to wider market forces
  - Fallback implementation schemes look expensive
- At the mercy of interconnect technology ‘futures market’

Scaling up For Phase 2

- All of the above... but much worse!
- Apply same principles to dense array and focal plane array solutions
- Its a scalability and affordability problem not a technical one
- Dense Array Tile processor for AA-Mid or PAF will look something like this ➔
- Need to find the right level of integration to remove cost and maintain adequate performance

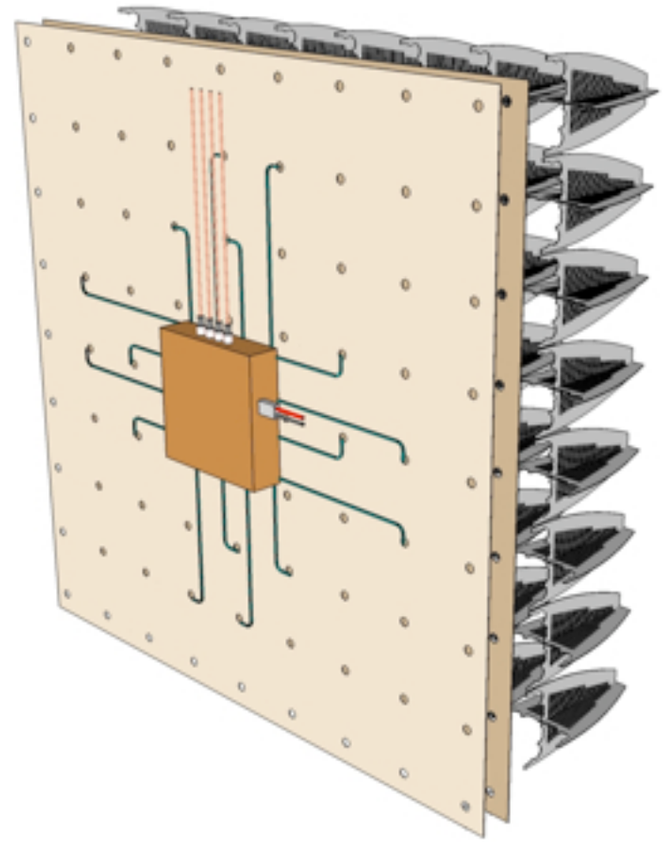

Multi Channel Integrated Digital Receiver Concept aka 'Tile Processor'

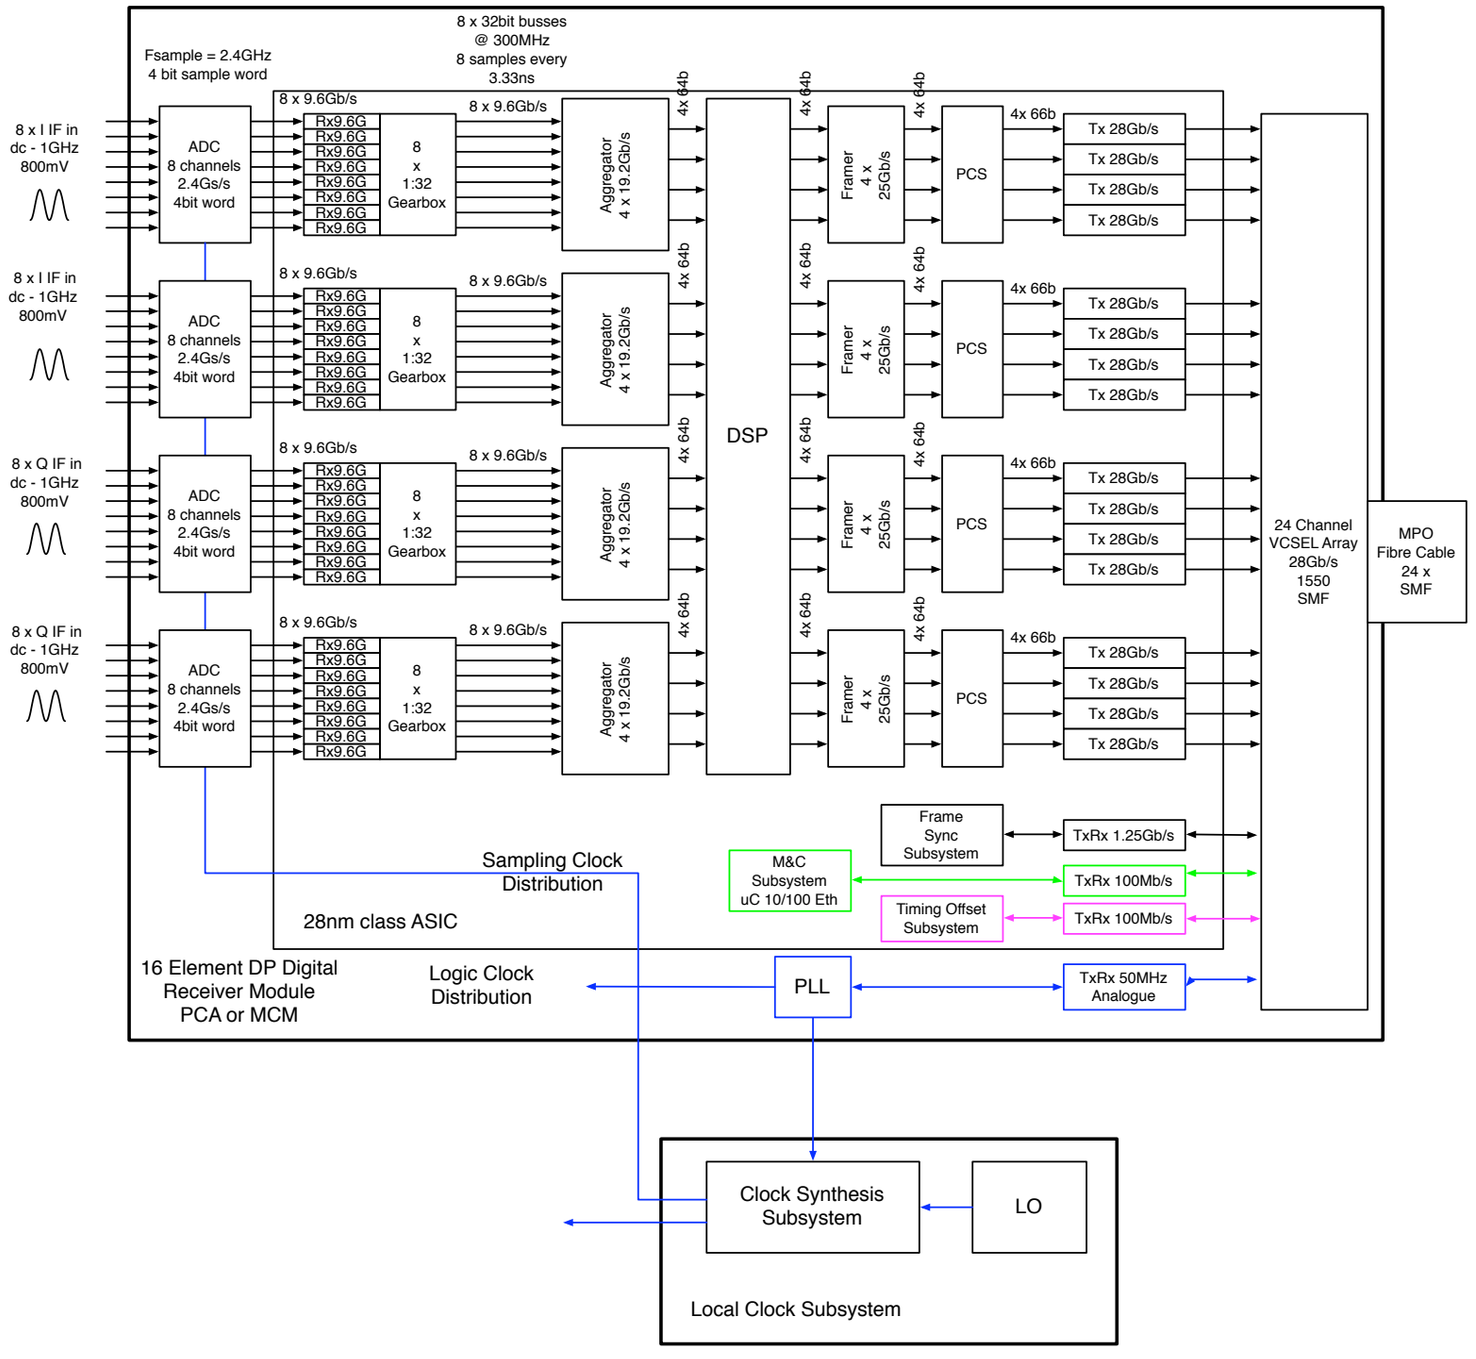

What Happens Next...

- Continue Design work on integrated solution based system reference design
- Architectural proving using 2-PAD and other Array based systems
- Update documentation
- Update the design concept with Station and Core processor Line Interfaces
  - Concepts understood and partially developed but additional work required
  - Costing of these interfaces
- Establish a real world reference design and a component level design database under way
  - We are way past the stage where generic costings are adequate
- Continue work on enclosure designs and RFI containment
- Explore cost implications of Analogue Transport to Integrated Tile Processor
  - Conventional
  - Analogue over Fibre
- Continue to build relationships with key technology suppliers and potential delivery partners;
  - Interconnect Vendors
    - current work is supported by TE (Formerly Tyco)
  - FPGA Vendors (Altera, Xilinx)

Conclusions

- We need to put constraints on the system performance in order to meet the cost targets
- Moving raw data samples over expensive links is expensive and unnecessary
  - This is not a datacomms network, we don't have to deliver unprocessed data
- We need to establish detailed reference designs for major options
- The scientists need to support these engineering activities by analysing the impact of the constrained design and buying into it...
  - Working at a theoretical level is now not adequate
- Can we build phase 1?
  - Yes, but not without compromises
- Can we build phase 2?
  - Yes, but we probably need to move to a high level of integration to make it affordable
  - We need to continue to develop real engineering relationships with technology partners