The SKA from a Communications Infrastructure Perspective

Charles Smith

Commonwealth Scientific and Industrial Research Organisation

Cnr., Vimiera and Pembroke Roads

Marsfield, New South Wales 2122

Australia

charles.smith@csiro.au

The SKA from a Communications Infrastructure Perspective

The SKA design process can be seen as an exercise in optimizing the transport and processing of very large volumes of data"

"It is desirable to be able to minimize design and construction costs by using commercially available equipment where possible, to exploit Moore's Law and available commercial products"

K. Van Der Schaff, R. Overeem "*COTS Correlator Platform*" Experimental Astronomy, vol 17, pp. 287-297, 2004

Topics of Discussion

A large component of the SKA will be the communications infrastructure required to interconnect the sensors and other elements of the system.

Data communications interface types are investigated for possible adoption in the SKA .

A resulting industry standards based, scalable data communications architecture for the SKA is then introduced.

Discussions here are based upon the published SPDO Memo 130 data rates for SKA Phase 1.

CSIRO

SPF Custom to COTS Interface

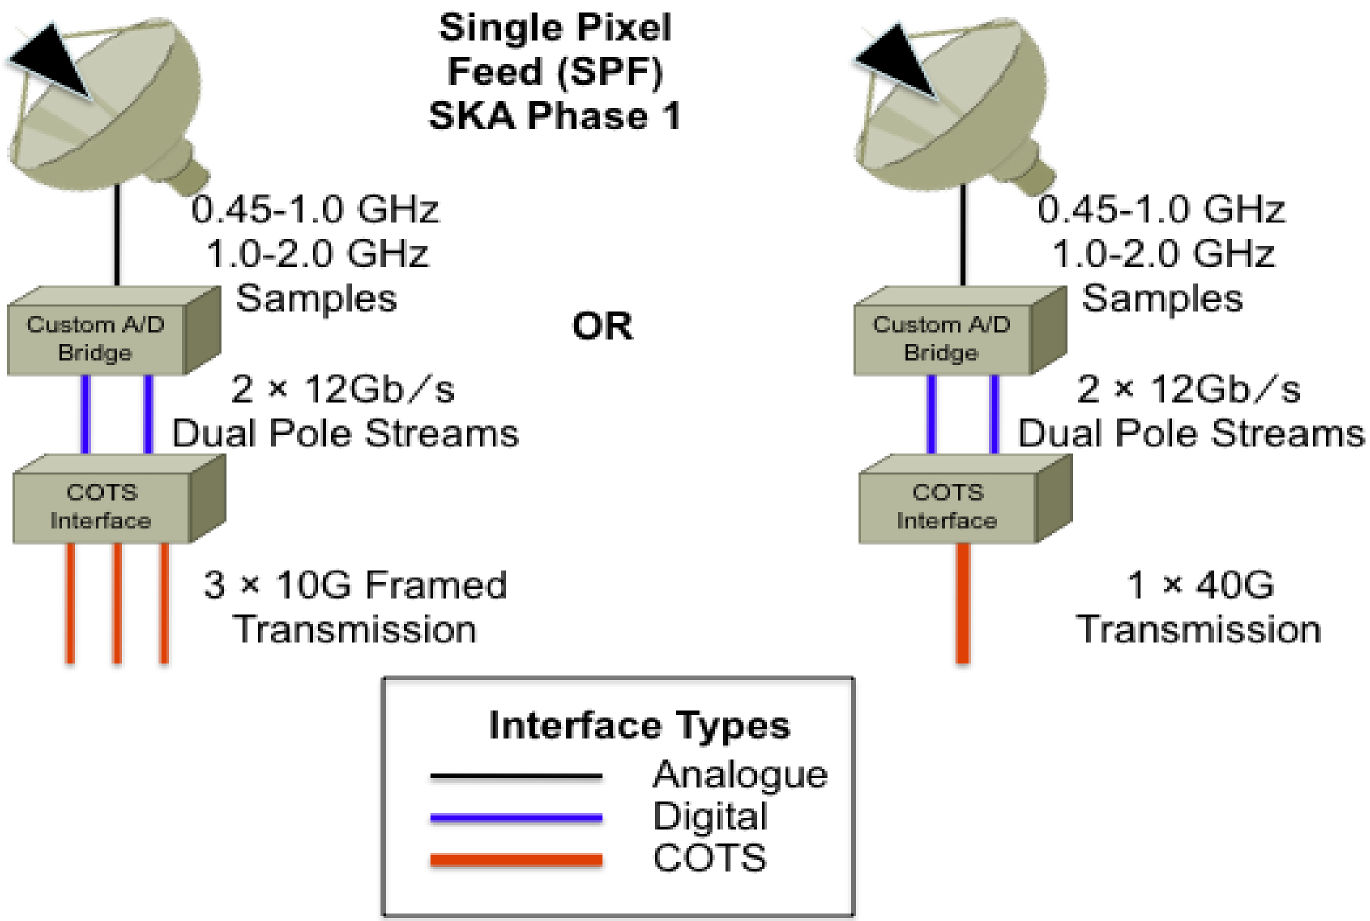

CSIRO

PAF Custom to COTS Interface

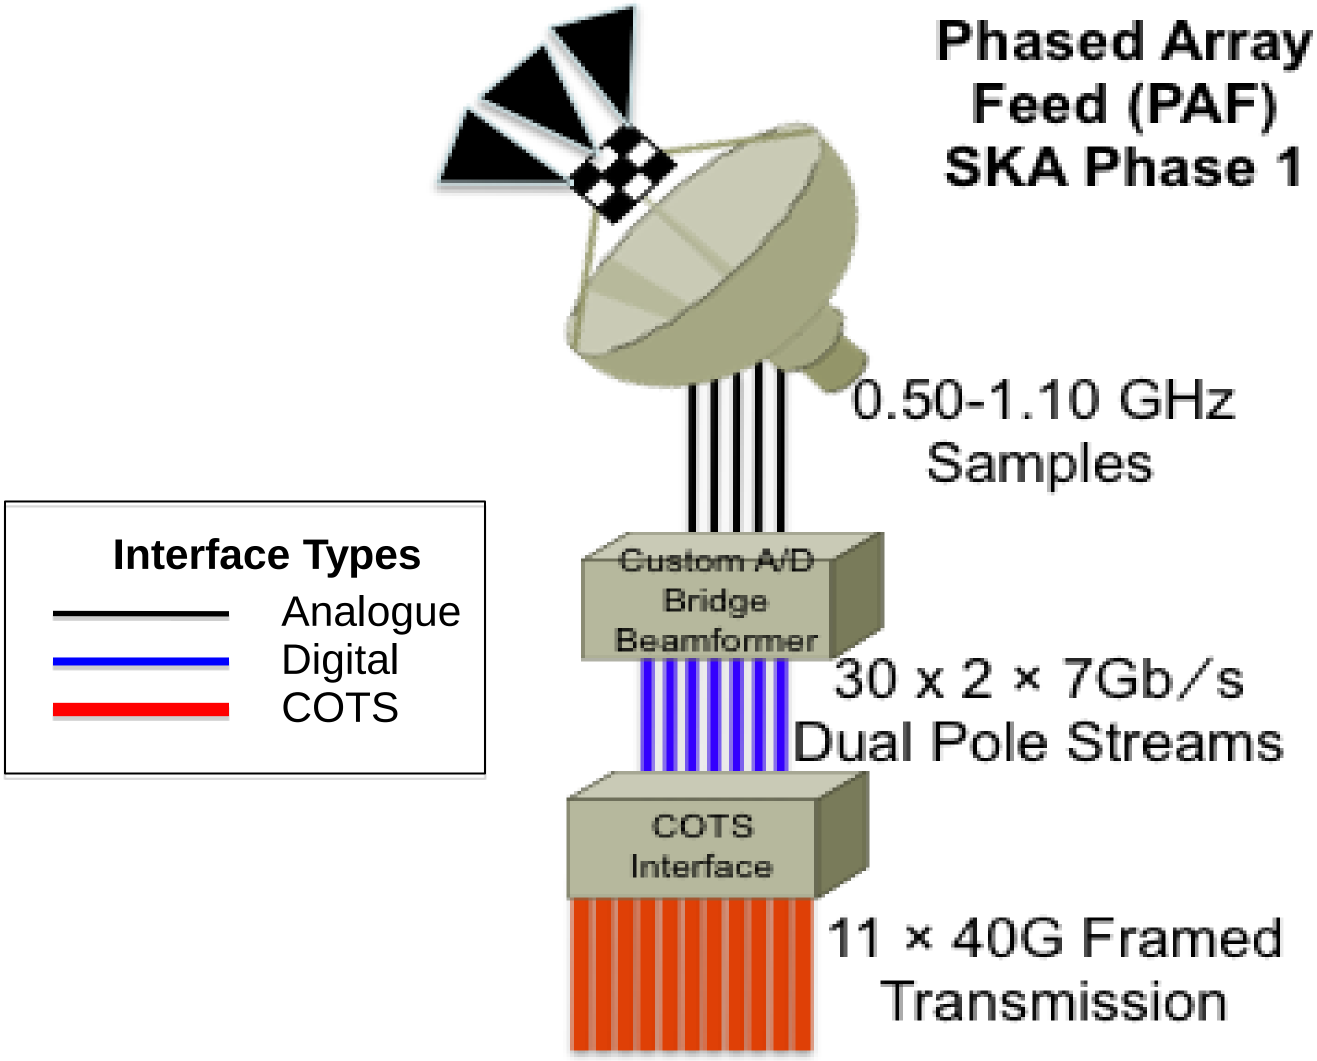

SAA Custom to COTS Interface

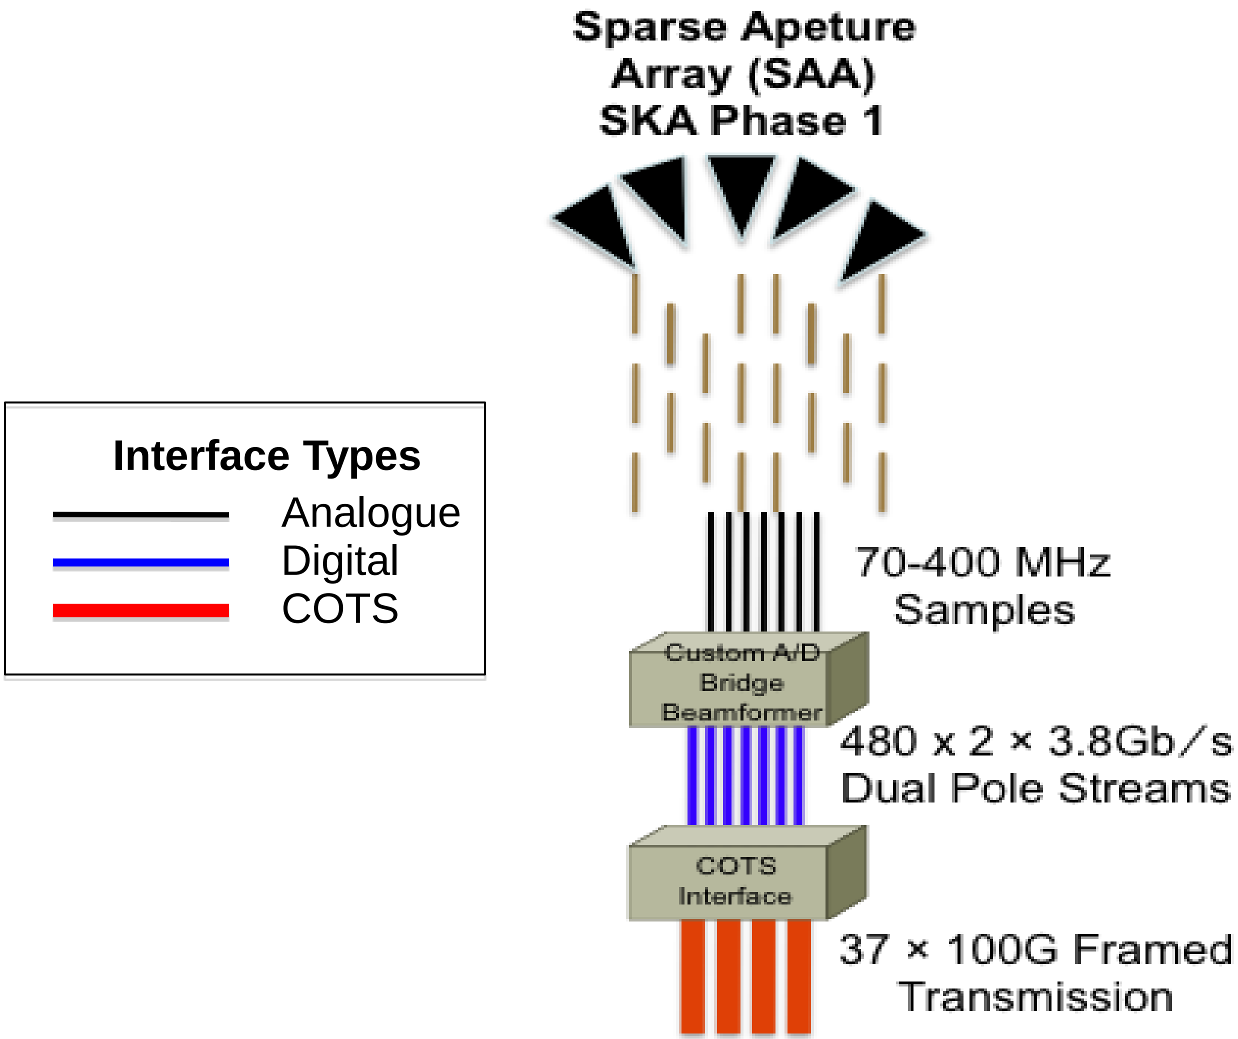

Choosing an Industry Standard Interface for Sensors

Output data rates across the differing sensor technologies vary widely.

Industry standards for digital data transmission that can satisfy the sampling requirement rates identified are SONET/SDH, Ethernet, and Fibre Channel.

Each have the capabilities to support the individual streams from sensors with interface addressing, multiplexing and aggregation to identify, encapsulate and reliably transmit the data.

They are supported in short haul high density Local Area Networks (LAN), Wide Area Network (WAN) aggregation interconnects and capable of long haul transmission via DWDM systems.

Infiniband Options & Capabilites

The Infiniband protocol and interface is currently limited to high-speed interconnections between nodes within close proximity and Data Centres.

It is a reasonable choice to investigate for use within the SKA correlation and supercomputing environments.

It is not an appropriate choice for the sensor to COTS interface given its distance limitations & current requirement to be subsequently bridged to other COTS technologies for wider transmission.

Development of Sensor to COTS Interfaces

For all sensor technologies, bridges between the streaming sensors and COTS transmission interfaces will need to be developed by the SKA community.

The 7Gbs^{-1} PAF beam stream and 3.8Gbs^{-1} SAA beam stream can be comfortably handled with individual COTS interfaces.

The 12Gbs^{-1} SPF data stream poses a complexity issue in that it exceeds 10G standards for base transport.

Moving up to the next industry standard of 40G would under-utilise the downstream transmission system.

Alternately the implementation of a buffering mechanism to support the aggregation of multiple 10G lower speed interfaces would address this issue.

SONET/SDH as a sensor to COTS Interface

Minimising the number of transmission COTS Interfaces can be achieved via statistical multiplexing in SONET/SDH.

A 3.8Gbs^{-1} SAA beam would require two 2.5G SONET/SDH circuits; the 1.2G additional overhead could not be regained unless a higher layer protocol is employed.

The higher layer protocol would need to provide a non-blocking buffering, aggregation and switching capability of multiple services onto a higher speed 10G or 40G uplink.

For a standards based solution a layer 3 routing capability is required and IP is the obvious choice.

The additonal cost involved in implementing a Layer 3 Intelligence is however the true restricting factor in selecting this method.

CSIRO

Ethernet & Fiber Channel as a sensor to COTS Interface

Ethernet and Fibre Channel non-blocking switching and aggregation provides standards based mechanisms using quality of service, inter-frame gaps and port buffering to maximise the average inbound frame rates to the aggregated higher speed 40G or 100G outbound interfaces.

As an example a 10GigE or 10G FC interface for a PAF 7Gbs^{-1} beam would allow for the 30% under-subscription capacity to be regained

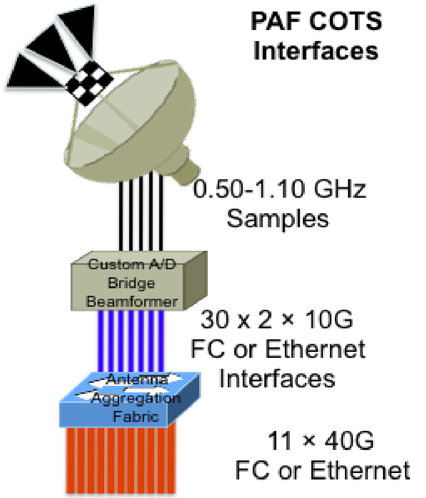

Sensor to COTS Interface for SPF

Ethernet and Fibre Channel non-blocking switching and aggregation provides standards based mechanisms using quality of service, inter-frame gaps and port buffering to maximise the average inbound frame rates to the aggregated higher speed 40G or 100G outbound interfaces.

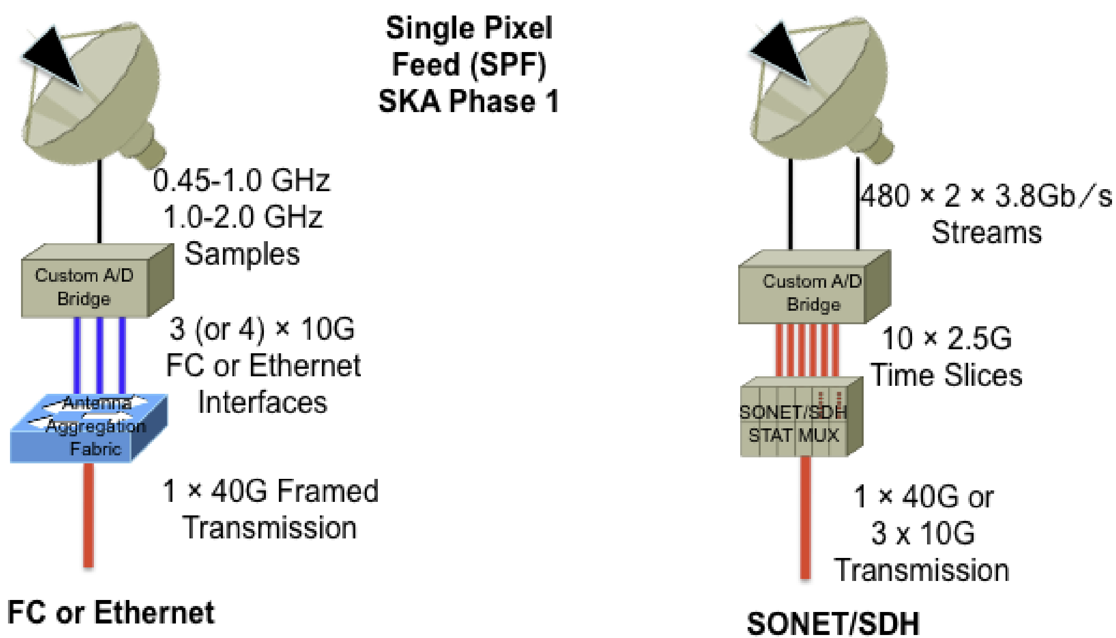

Sensor to COTS Interface for PAF

10GbE or 10G FC as a standards based interface for a PAF 7Gbs^{-1} beam allows for the 30% under-subscription capacity to be regained in a switching fabric if Ethernet or Fibre Channel is chosen.

If SONET/SDH is used then the 7Gb could be aggregated most efficiently as 3 x 2.5Gbs^{-1} time slices with an aggregation loss overhead of 0.5Gbs^{-1} for each PAF beam.

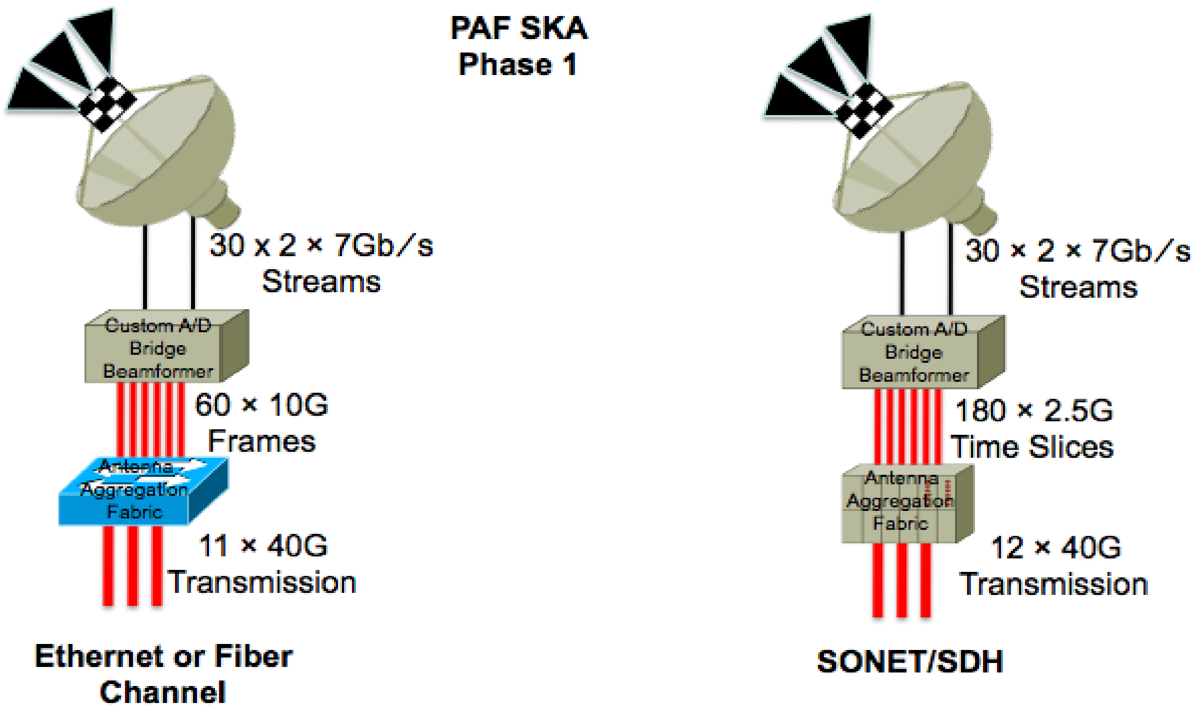

Sensor to COTS Interface for SAA

An SAA station with 3.8Gbs^{-1} from each beam-formed pole allows for a 10GigE or 10GFC interface to comfortably supply data with the additional 20% under-subscription capacity regained in the port rate limiting, quality of service, switching and aggregation capabilities.

SONET/SDH poses an inefficient aggregation solution to be adopted as either 4 x 2.5G time slices or 1 x 10G interface is fully utilised for the single sampling function resulting in an 2.4G waste.

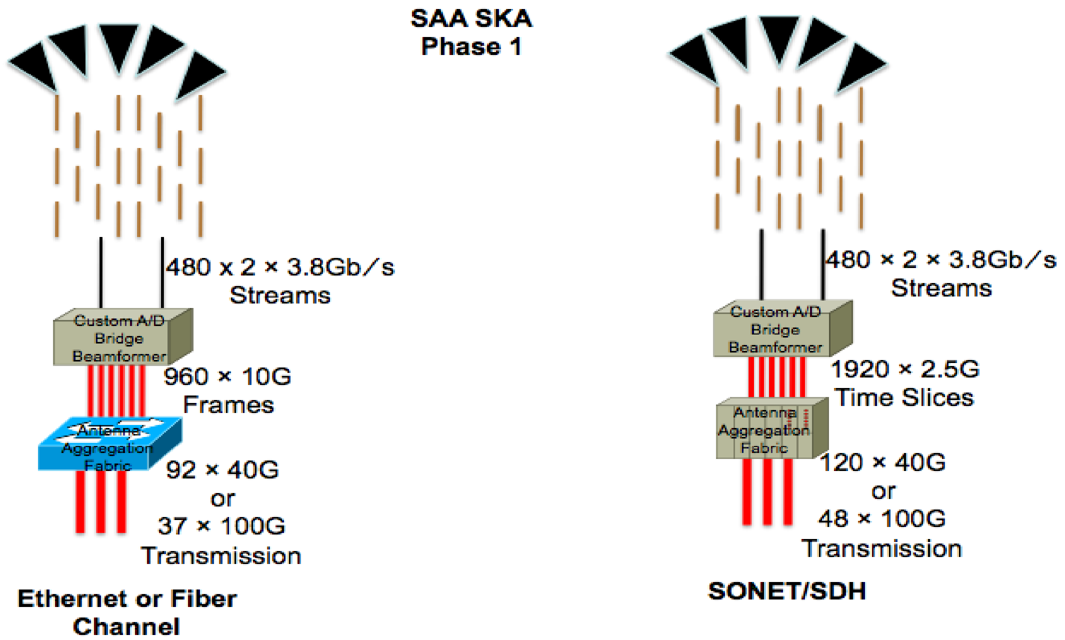

CSIRO

Fiber Channel as a SKA Sensor COTS Interface

Fibre Channel suits simple aggregation functions but needs to be planned and expanded very carefully as the native protocol is blocking based.

The standard demands a transmitted frame be delivered and will block other traffic until it has been delivered.

This does guarantee delivery of every frame however if there is an imbalance in the system then the link will block.

The buffer credit mechanism is used to offset this issue, however it requires careful design of each link in the network and which varies based upon distance and propagation.

CSIRO

Ethernet as a SKA Sensor COTS Interface

Ethernet provides for a common Layer 2 interface across all sensor types.

It is also a common interface between sensors and the beamforming / correlation compute tasks.

Expected growth in sensor rates for SKA Phases 2 and 3 are catered for with the introduction of 40GigE and 100GigE while allowing SKA Phase 1 systems to co-exist and grow.

It allows for the upstream optimization of traffic throughput and flexible growth by integrating individual interface under-subscriptions by utilizing switching aggregation fabrics, port buffering/queing & QoS mechanisms.

Ethernet as a SKA Sensor COTS Interface

Server interconnections are moving to native support for 10GigE and are becoming available at 40GigE by 2014 and 100GigE by 2020.

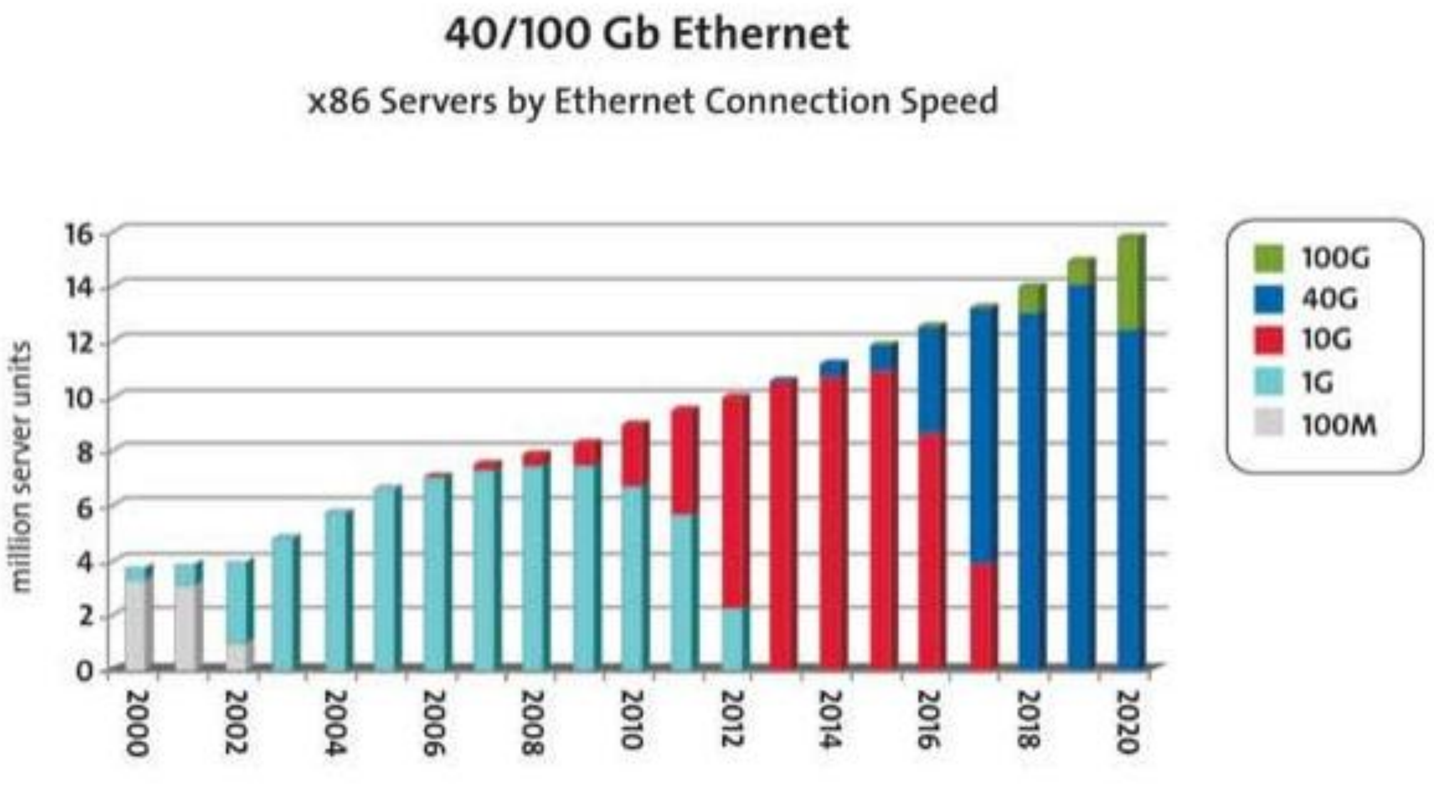

Source: Intel/Broadcom (Nov, 2007)

Ethernet as a SKA Sensor COTS Interface

Accelerated annual port sales forecast vs. the drop in average sales price over the next 4 years show 10GigE to be a cost effective technology to base SKA Phase 1 A-D sensor outputs as the COTS interface of choice

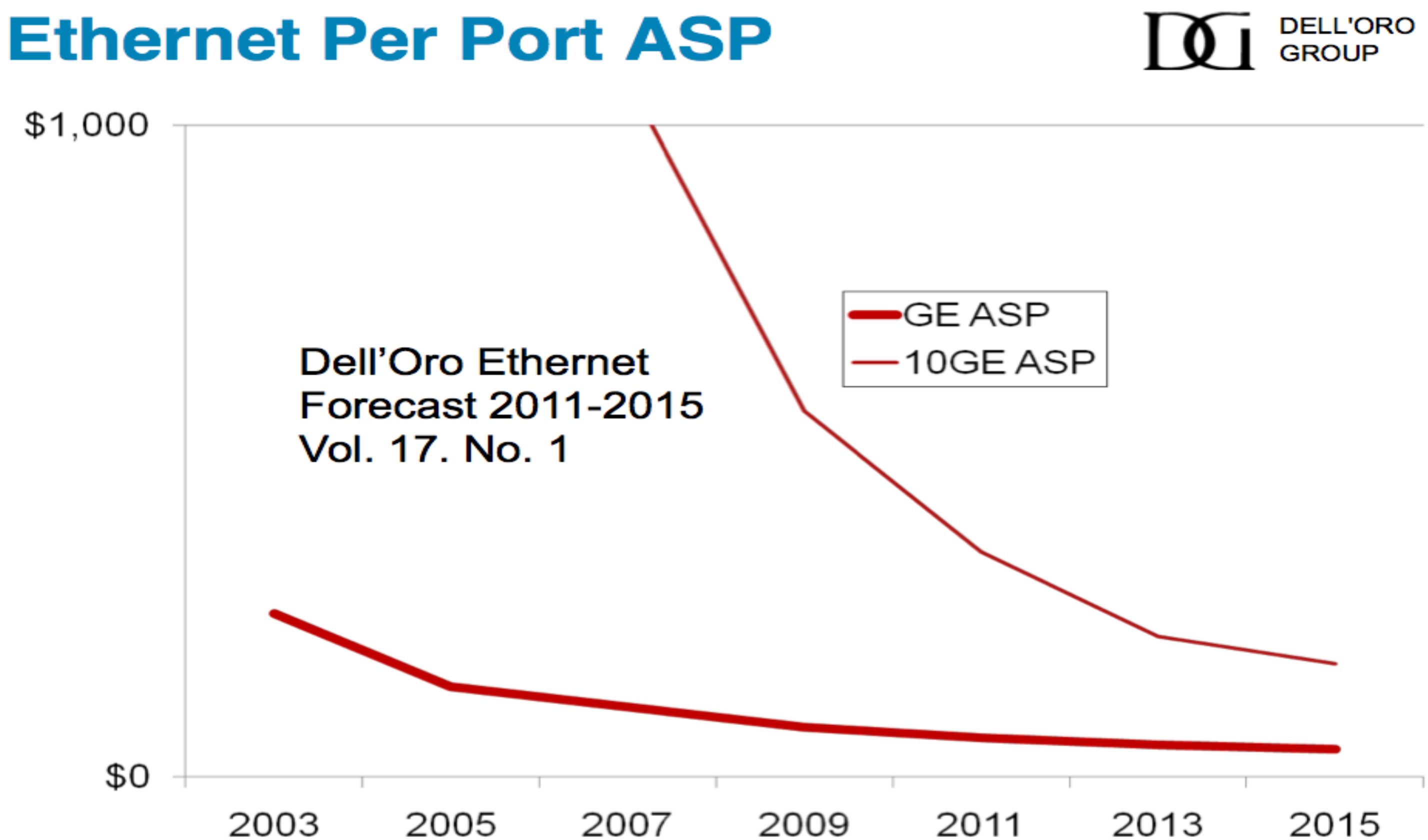

Ethernet as a SKA Sensor COTS Interface

The defined standard of 40GigE & 100GigE, (IEEE Std 802.3ba-2010) allows for the smooth transition and upgrade path of today's networks.

Ethernet standards are very well understood across the communications industry, the scientific community, the engineering sectors and business sectors as it has proven price/performance and a predicted growth pattern.

CSIRO

SKA Phase I SPF Data Flow & Aggregation

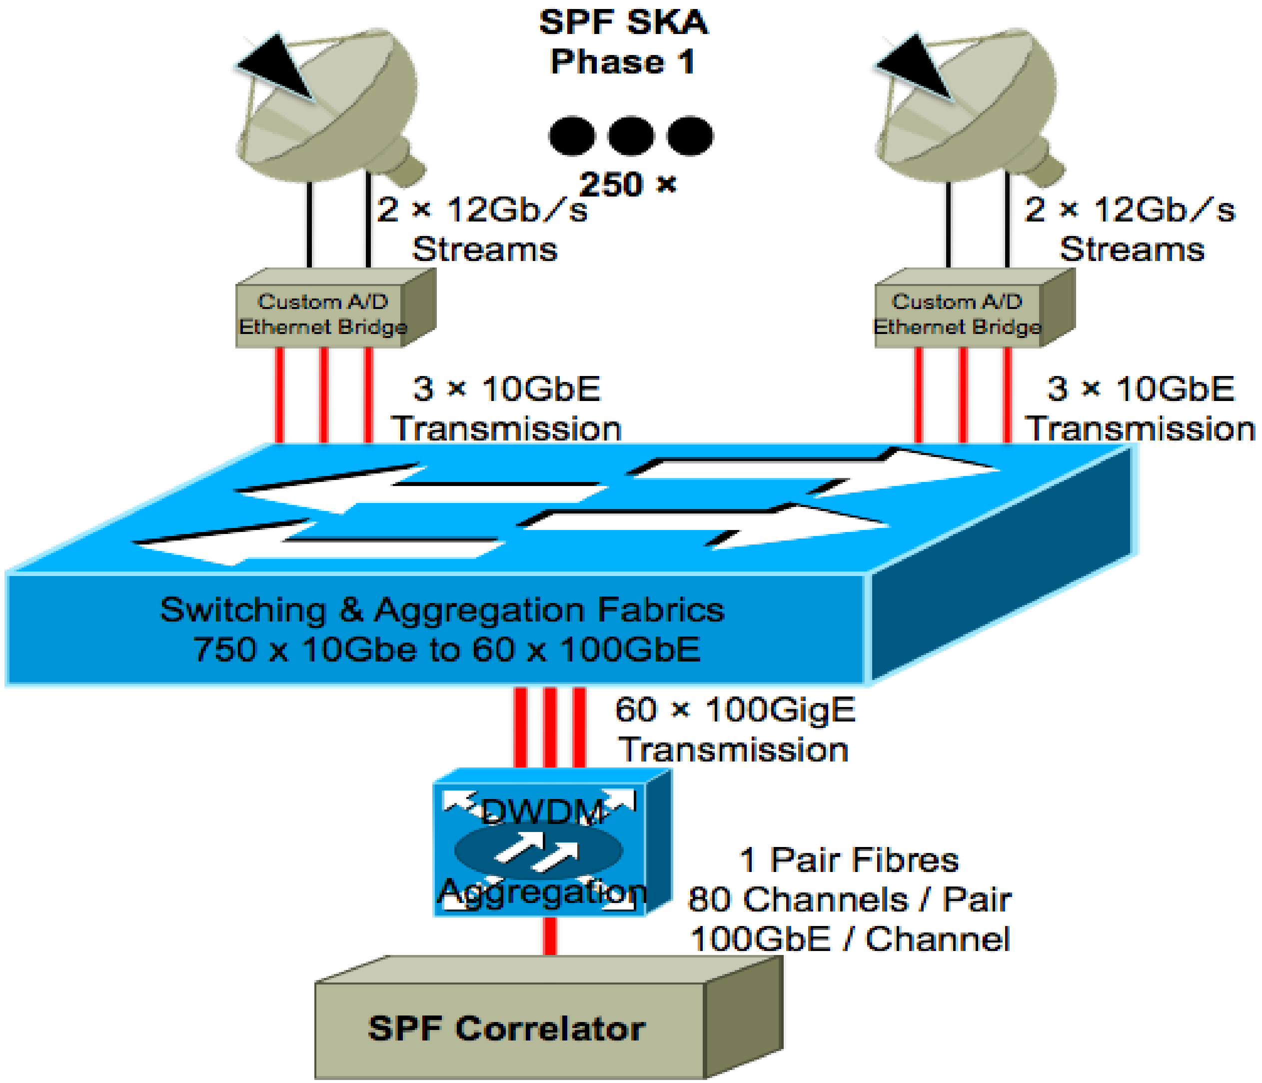

SKA Phase I PAF Data Flow & Aggregation

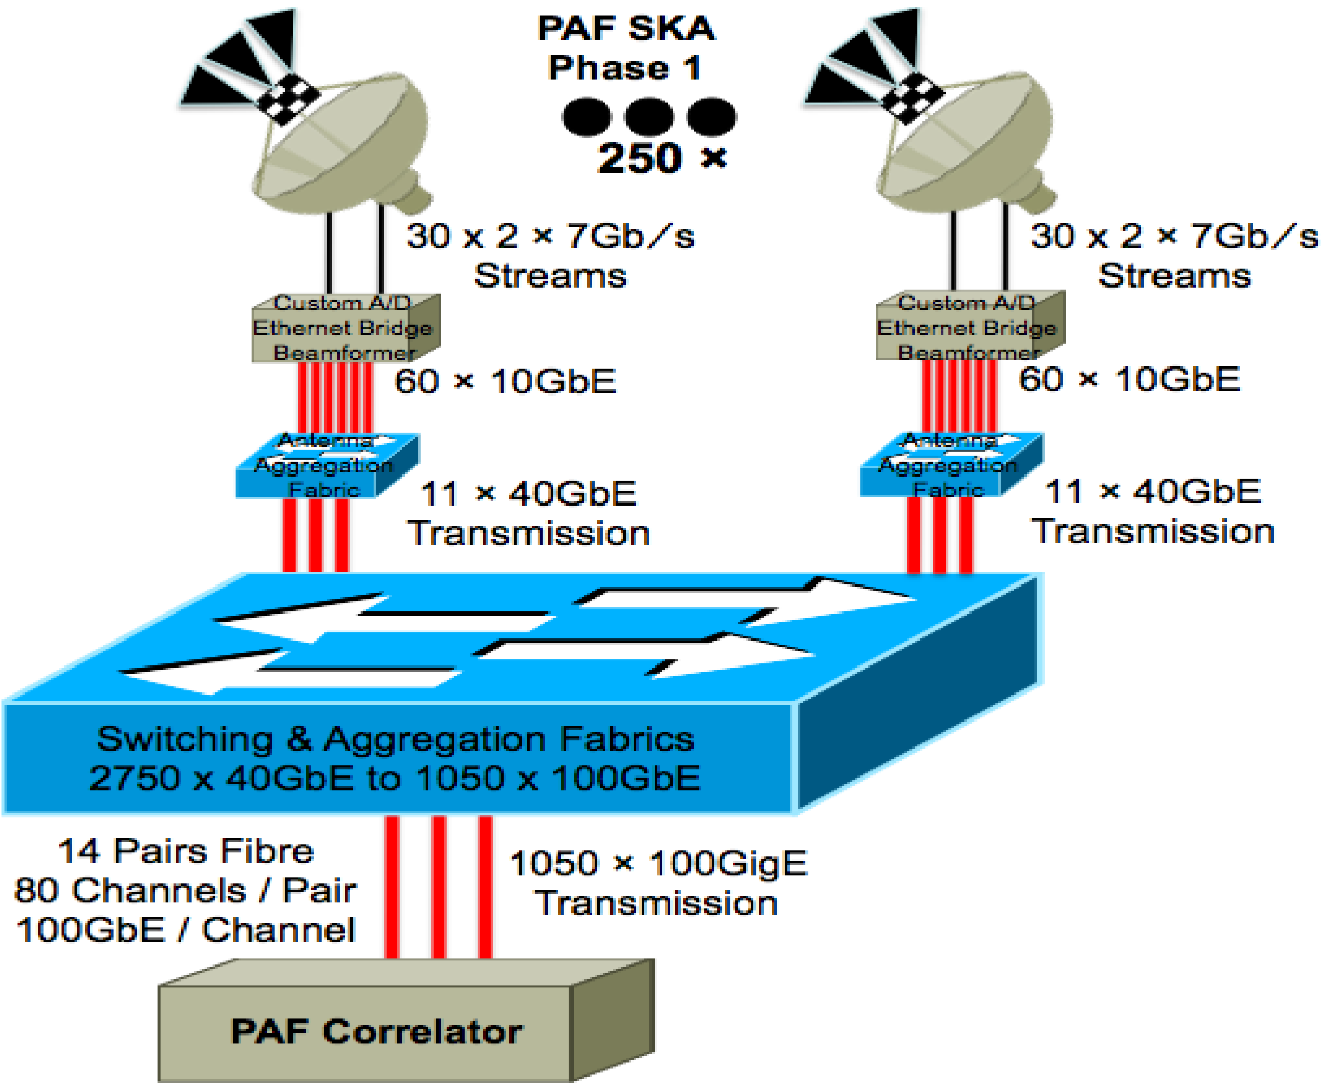

CSIRO

SAA Phase I PAF Data Flow & Aggregation

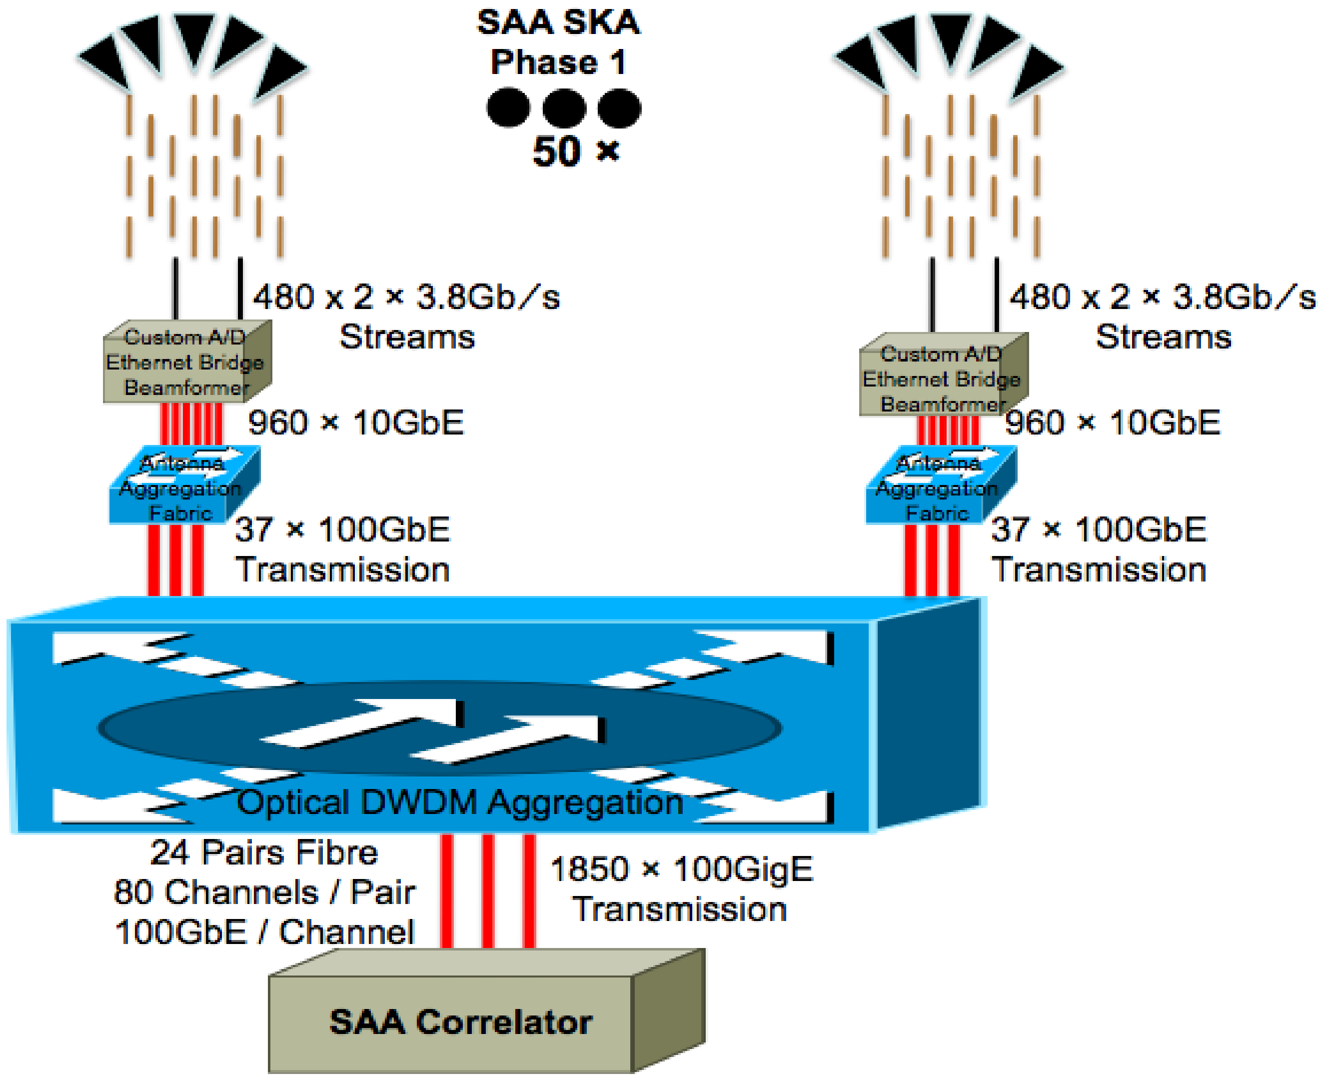

CSIRO

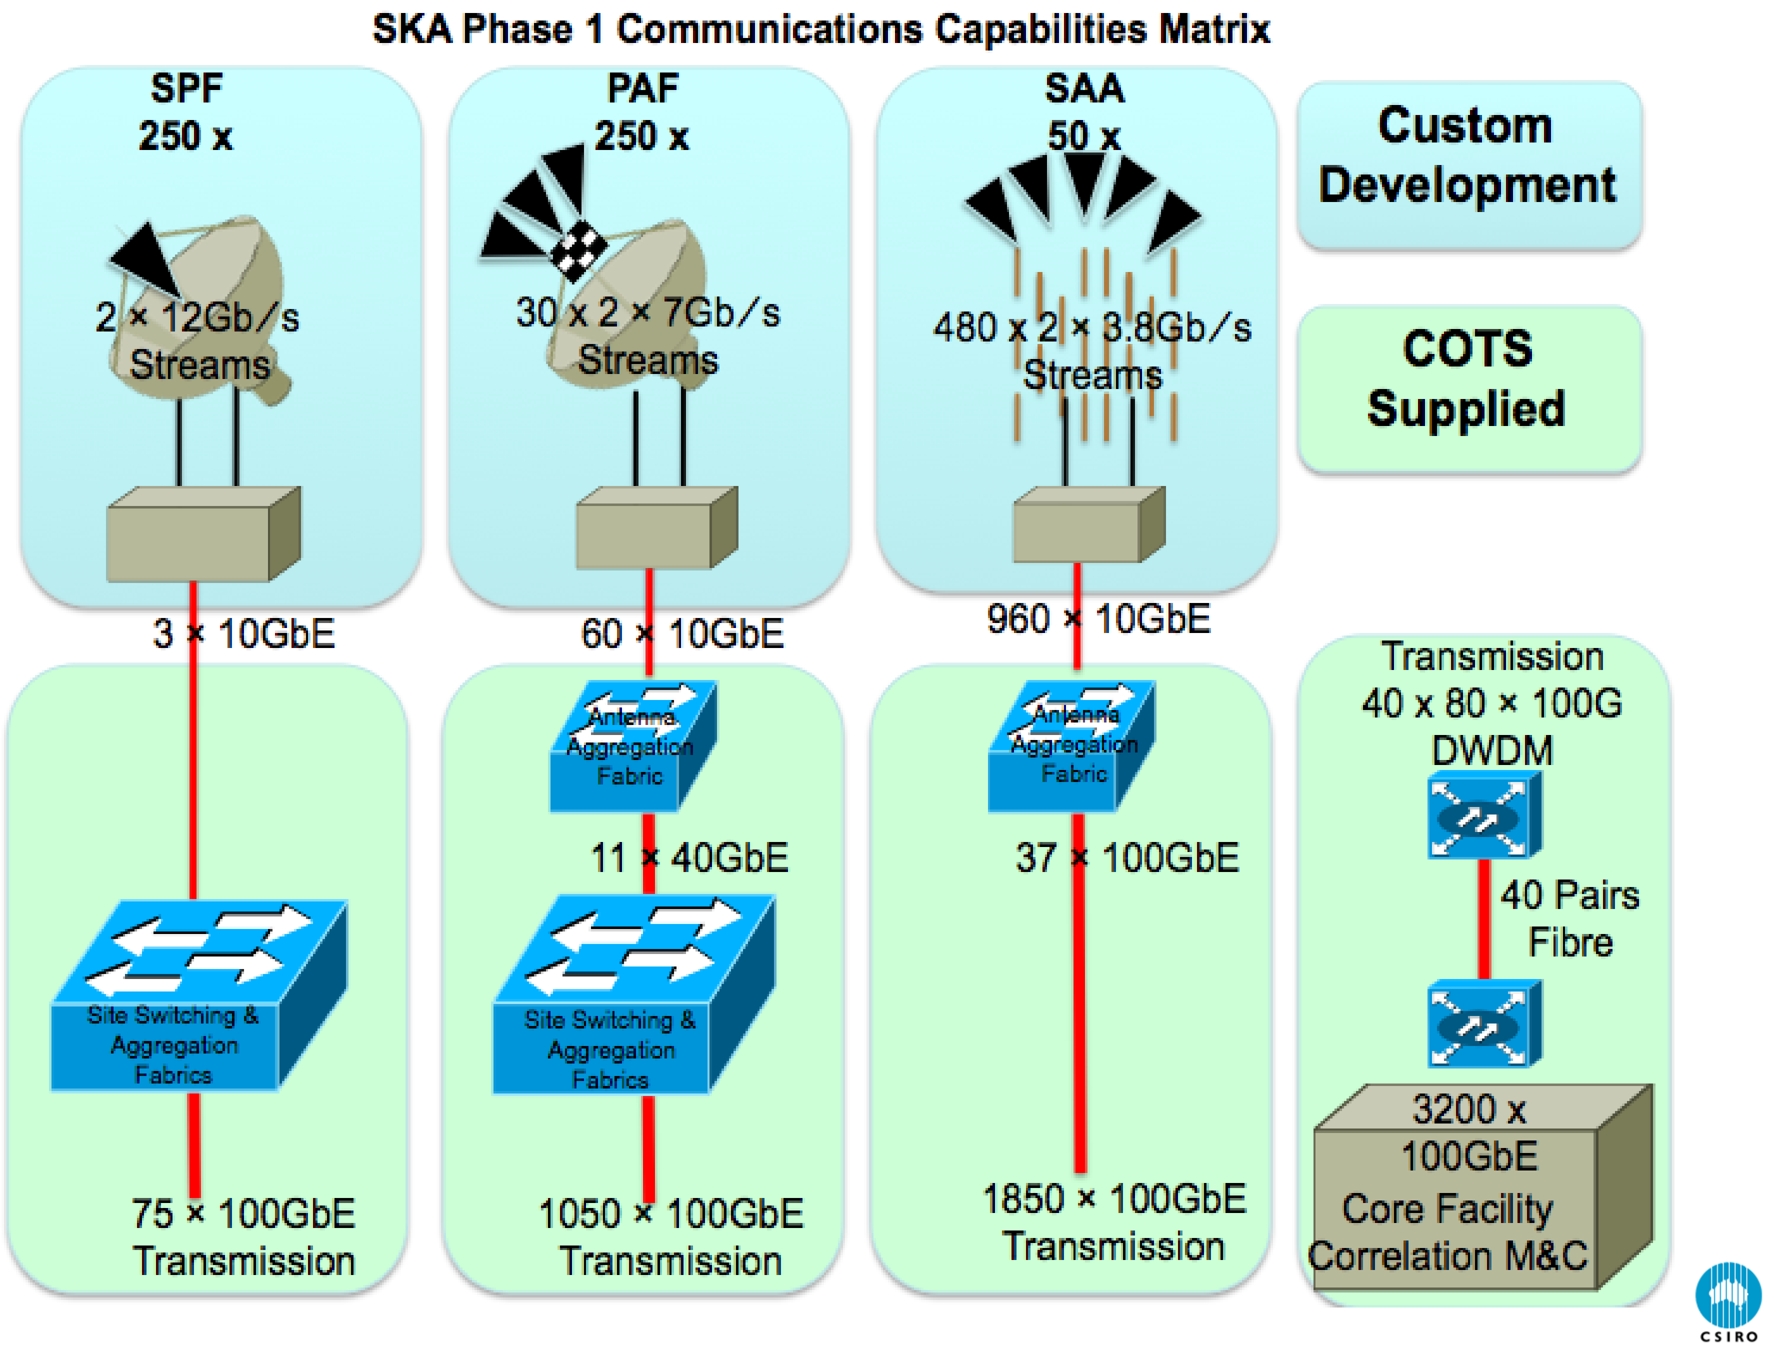
SKA Phase 1 Communications Capabilities Matrix
SPF
250 x
2 × 12Gb/s
Streams
3 × 10GbE
Site Switching &
Aggregation
Fabrics
75 × 100GbE
Transmission
PAF
250 x
30 x 2 × 7Gb/s
Streams
60 × 10GbE
Antenna
Aggregation
Fabric
11 × 40GbE
Site Switching &
Aggregation
Fabrics
1050 × 100GbE
Transmission
SAA
50 x
480 x 2 × 3.8Gb/s
Streams
960 × 10GbE
Antenna
Aggregation
Fabric
37 × 100GbE
1850 × 100GbE
Transmission
Custom
Development
COTS
Supplied
Transmission
40 x 80 × 100G
DWDM
40 Pairs
Fibre
3200 x
100GbE
Core Facility
Correlation M&C
CSIRO

Summary

Industry has switching and transmission systems in-place that will satisfy the SKA raw data communications requirements for Phase 1 of the project today and standards are being ratified to satisfy the requirements for Phases 2 and 3.

Intelligence in bandwidth aggregation capabilities coupled with systems growth support is a critical factor for building a scalable SKA.

Using a standards based well-understood communications interface allows for the close liason between between industry, the astronomy researchers and systems engineering researchers architecting, building and subsequently operating the SKA.

CSIRO