Current Developments in the NRC Correlator Program

Lewis B.G. Knee (on behalf of the NRC Herzberg Correlator Group)

ALMA Developers Workshop
Chalmers University of Technology, Gothenburg
May 26, 2016

Main NRC Activities in SKA Development

CSP: Central Signal Processing

- Consortium led by NRC with MDA corporation
- Participation from CSIRO, Manchester, Swinburne, many others
- SKA1-Mid.CBF leads at NRC are Brent Carlson and Michael Rupen

SDP: Science Data Processing

- Consortium led by Cambridge with ASTRON, ICRAR, SKA-SA, NRC
- Science data products delivery (Severin Gaudet, CADC)

DSH: Dish

- Consortium led by CSIRO with SKA-SA, China, NRC, Sweden, Italy, others
- SPF cryogenic LNAs (Frank Jiang)
- SPF digitizers (Kris Caputa)

AIP: Advanced Instrumentation Program

- Band 1 - 2 (0.7 - 1.5 GHz) uncooled PAF (Bruce Veidt)
- Band 4 (2.8 - 5.2 GHz) cryogenic PAF (Lisa Locke)

SKA1-Mid.CBF General Specifications

SKA1-Mid will have 197 antennas, 150 km array extent

Correlator

- 5 GHz/polarization
- 64 K channels/polarization/baseline
- 4 x 16 K channels/polarization/baseline zoom mode

Beamformer

- PSS (pulsar search) up to 1500 beams with 300 MHz bandwidth
- PST (pulsar timing) up to 16 beams with 2.5 GHz/polarization (total data input rate 80 Gb/s per beam), ~ 200 ns impulse response for high time resolution pulsar timing

VLBI mode

- beamform 4 full bandwidth beams, each beam formed from all antennas
- VLBI standard rates

16 completely independent sub-arrays

10 petaflops for combined CBF

Correlator/Beamformer Block Diagram

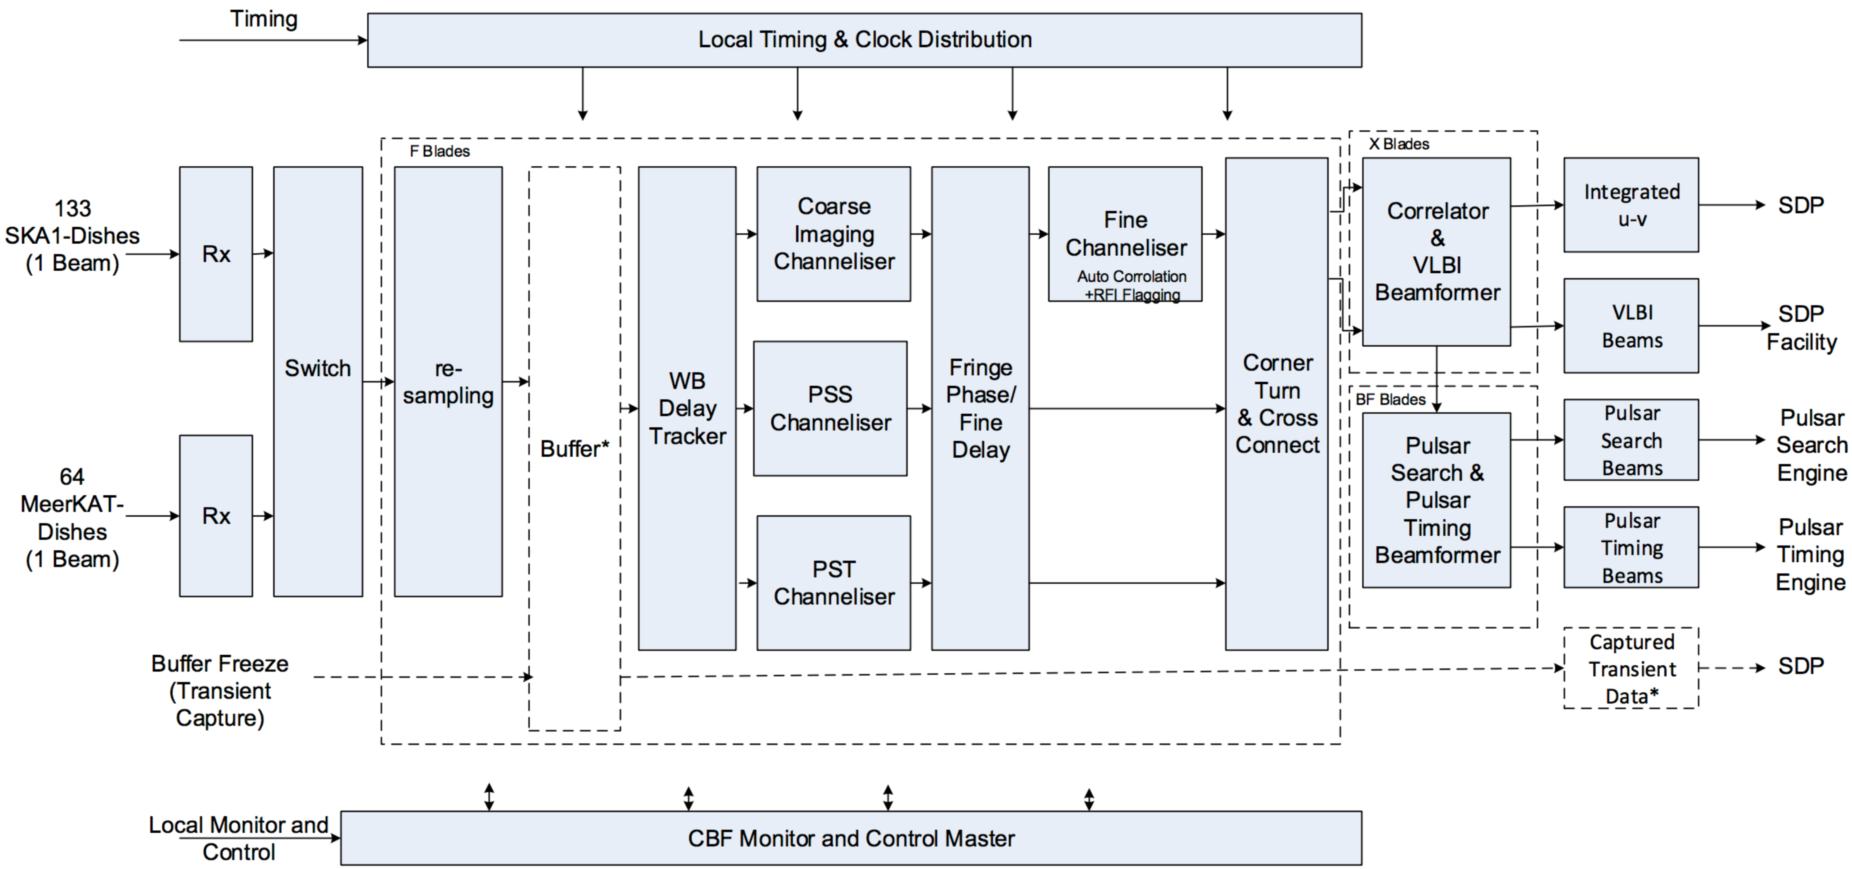

Note that Fringe Phase/Fine Delay has been moved to the Resampler for higher precision and signal fidelity

NRC Approach to SKA1-Mid.CBF Development

Building on the PowerMX framework

Set of layered specifications defining

- Motherboard and mezzanine card board form factors
- Connector pinouts
- Power supply, control, and monitor requirements

PowerMX provides

- Integrated and flexible platform for computing
- Flexible platform for connection of processing devices via high-speed SERDES
- High density and high I/O count

Hopefully can cost-effectively keep up with the technology curve, while delivering high-performance robust systems

NRC Approach to SKA1-Mid.CBF Development

Other advantages of the PowerMX approach

- Standard specification could be adopted by industry (buy rather than build)
- Allow other groups to make contributions in defined areas under the umbrella of a well-defined framework
- Permits our “custom” technology development to stay relevant/ competitive in the COTS environment
- Launching pad for other large scale computing system projects (e.g. SKA2, ngVLA, ALMA?)

- Working group*: NRC and AUT (New Zealand)

* Initially also included ASTRON, SKA-SA, STFC (ATC and RAL), Manchester, AASL (UK), Oxford

PowerMX in a Nutshell

- Motherboard with processor (“PMXM”) and I/O (“PMX_IOC”) mezzanine card sites
- Motherboard provides services to mezzanine cards
  - LVDC power, power sequencing
  - 1 Gb/s (and/or 10 Gb/s) Ethernet switched network using SGMII-CDR
  - Voltage and temperature monitoring
  - Additional monitor and control lines: I2C, 1-Wire, SMI
  - SERDES connectivity: 384 pairs to 4-site motherboard; 28 Gb/s per SERDES (total aggregate 10 Tb/s I/O possible)
- Each PMXM site (4 per motherboard) is the same
- Each PMX_IOC site (8 per motherboard) is the same and accepts monolithic or partial mezzanine cards
- All specifications on www.powermx.org

PowerMX Motherboard

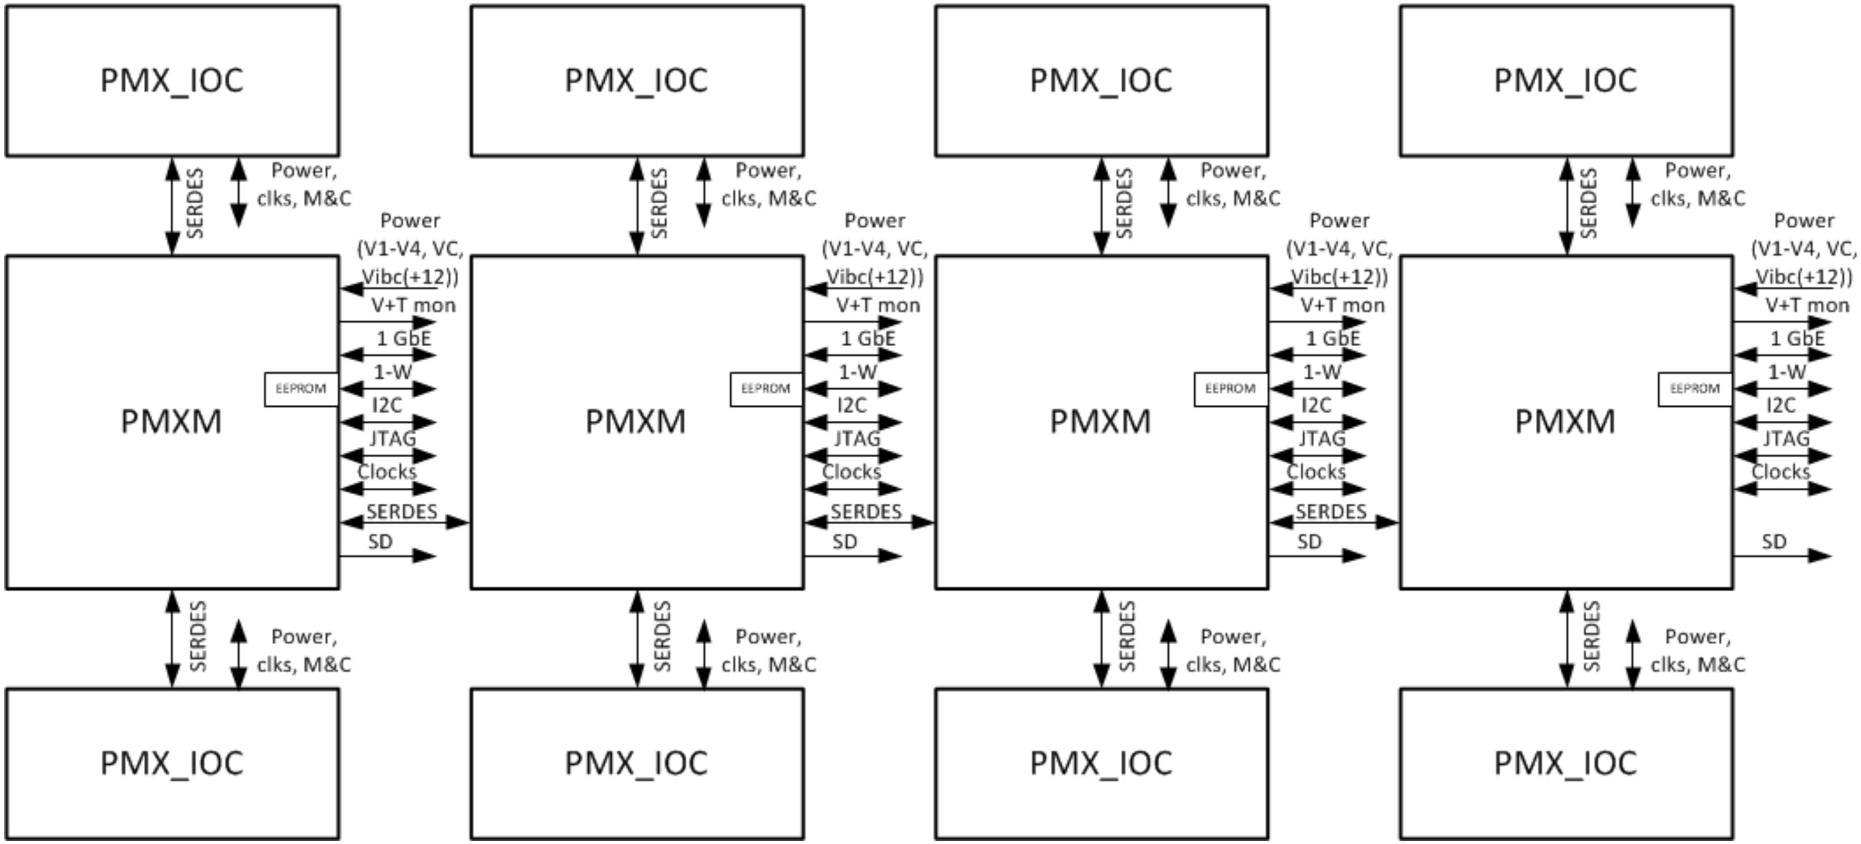

PowerMX Specifications Layers

PMX.1: Base Specification

- Foundational specifications for motherboard, mezzanine cards, connectors, pins, monitor and control, principles of other layers
- Packaging/enclosure not defined (e.g., liquid or air cooled 19” rack-mount server boxes)

PMX.1.1: Defines an alternative PMXM-PMXM SERDES mesh connection good for Fourier domain BFs and correlators

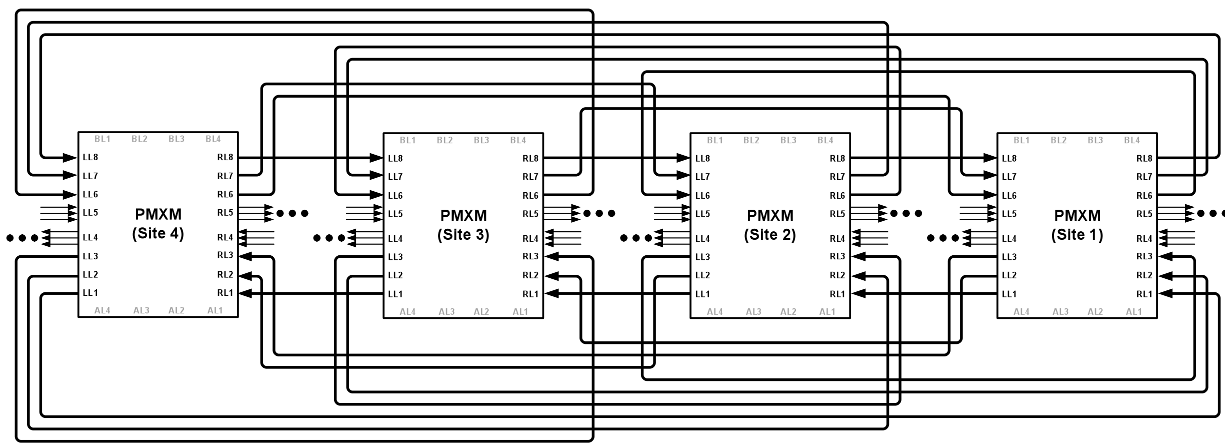

PMX.1.2: I/O card dimensions, bezels, panel cutouts, etc

PMX.1.3: Classes/types of PMX12 I/O pin assignments for compatibility

PowerMX Specifications Layers

PMX.2: Supervisory Monitor and Control (SMC)

- Defines standard methods (UDP/IP/AXIoE) for talking to motherboard SMC controller for providing/controlling supervisory services to mezzanine cards
- Allows for implementation-specific services/monitor points the motherboard provides
- SMC client(s) can be anywhere on the network
- SMC controller queries mezzanine cards to find out voltages, power, communications SERDES, etc. Must be compatible before it will power them up. Up to application to ensure applications are compatible

PMX.3: Application Monitor and Control (AMC)

- Defines standard method to query application (mezzanine card) as to what services/protocols it supports
- AXIoE suggested to be one for FPGA register read/write via host, otherwise communication protocols not constrained by the specification (heterogeneous computing)
- AXIoE provides uniform building-block approach in firmware and software
- Multi-client support

SKA1-Mid.CBF

Mid.CBF is the focus of our efforts and a good test case for PowerMX

- Flexibility
- Performance
- Power: 14 nm FPGA power in the 200 - 300 W/chip range. HERON motherboard should be able to deliver 288 A per PMXM site

So far, the flexibility inherent in the PowerMX framework has been very helpful in not having to say “NO” to changing/ increasing requirements

HERON Motherboard

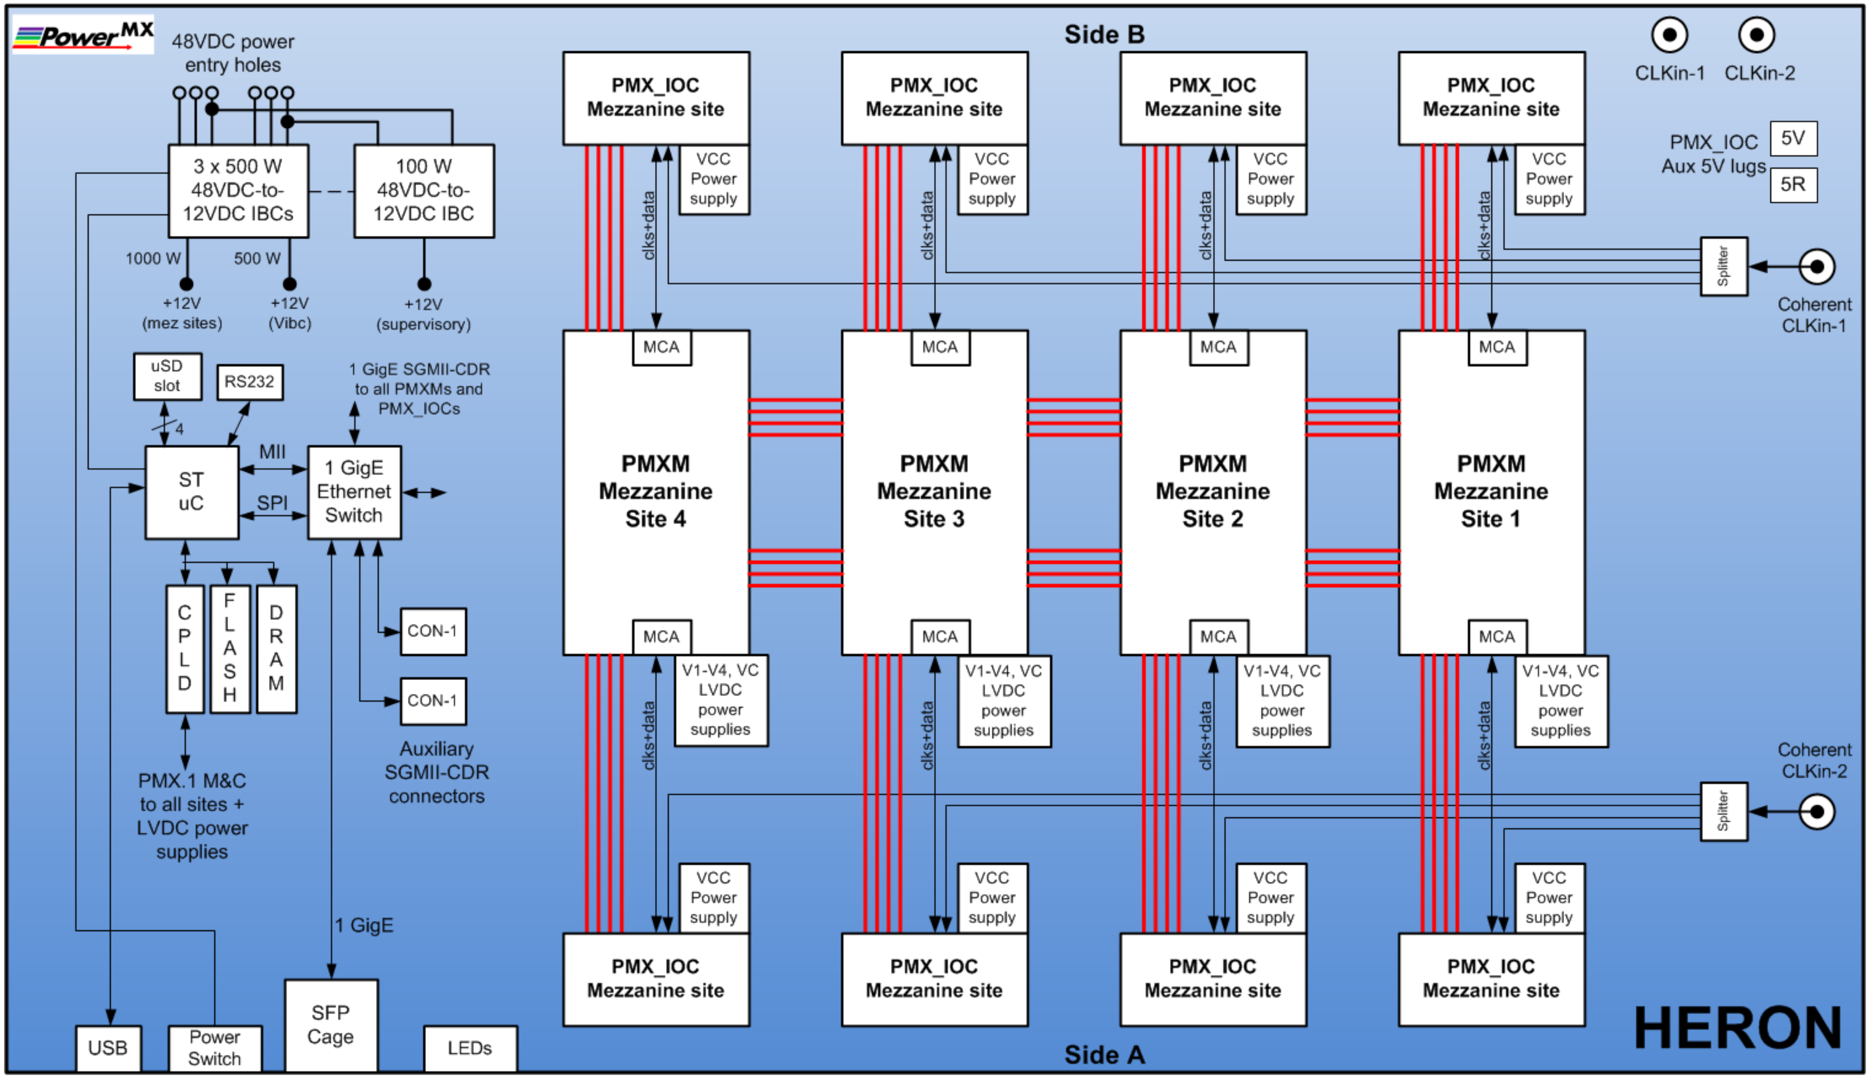

HERON Motherboard and HUMMINGBIRD-T I/O Card

About 300 HERON motherboards for SKA1-Mid.CBF

Board size: 17.244” x 8.268”

Board thickness: 125 mils

Board material: Megtron-6

Layers: 28

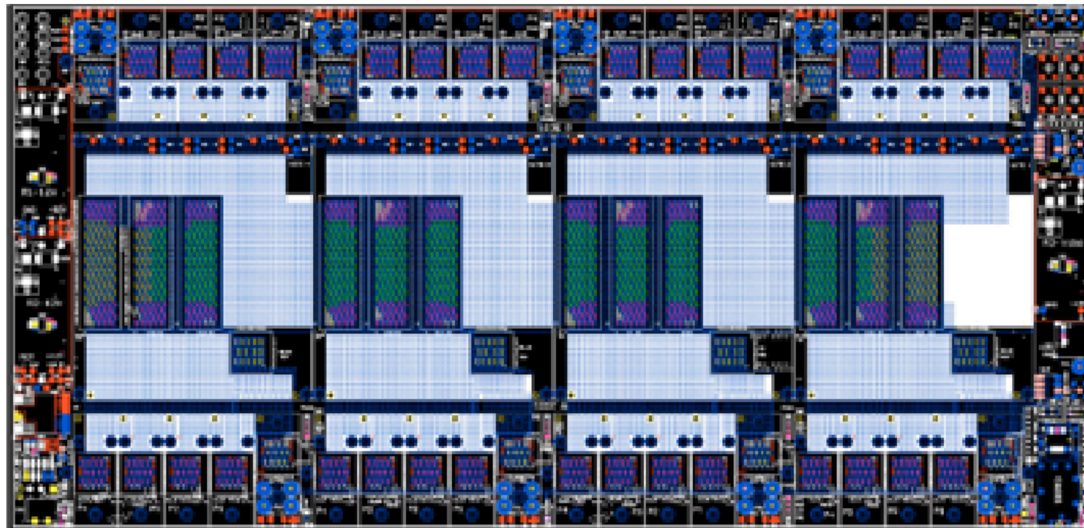

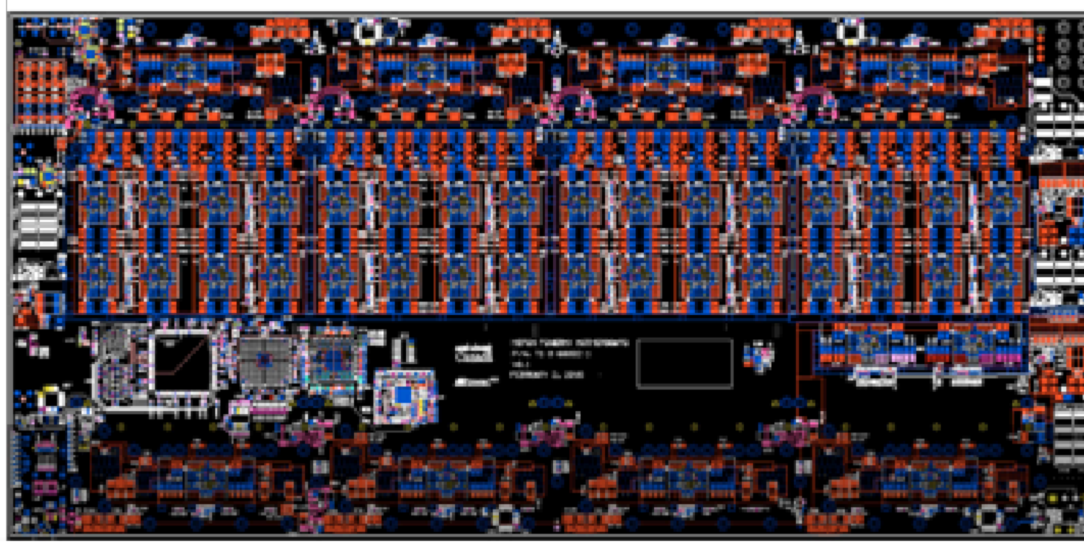

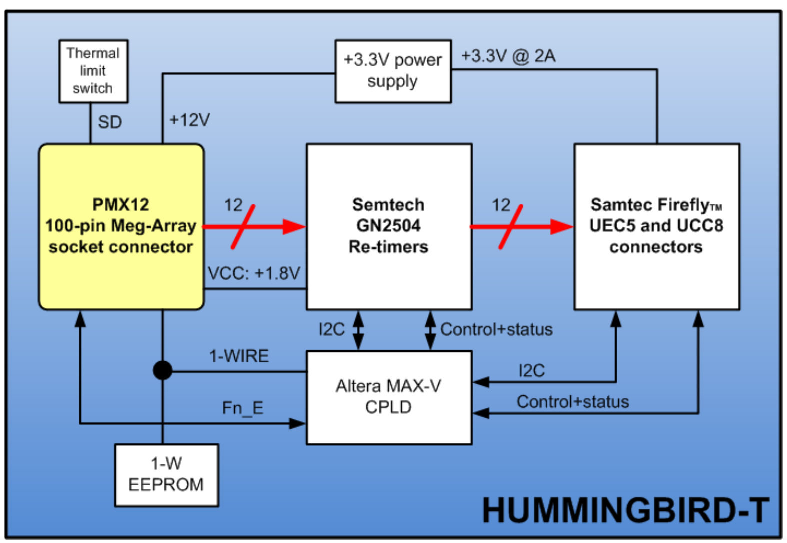

Correlator Mezzanine Cards

First 32-processor correlator card P32S-A10 (Arria-10 FPGA) being built now for testing this year. Upgraded RAVEN card (Stratix-10 FPGA) with 5 - 8 times the horsepower and better power/performance under development

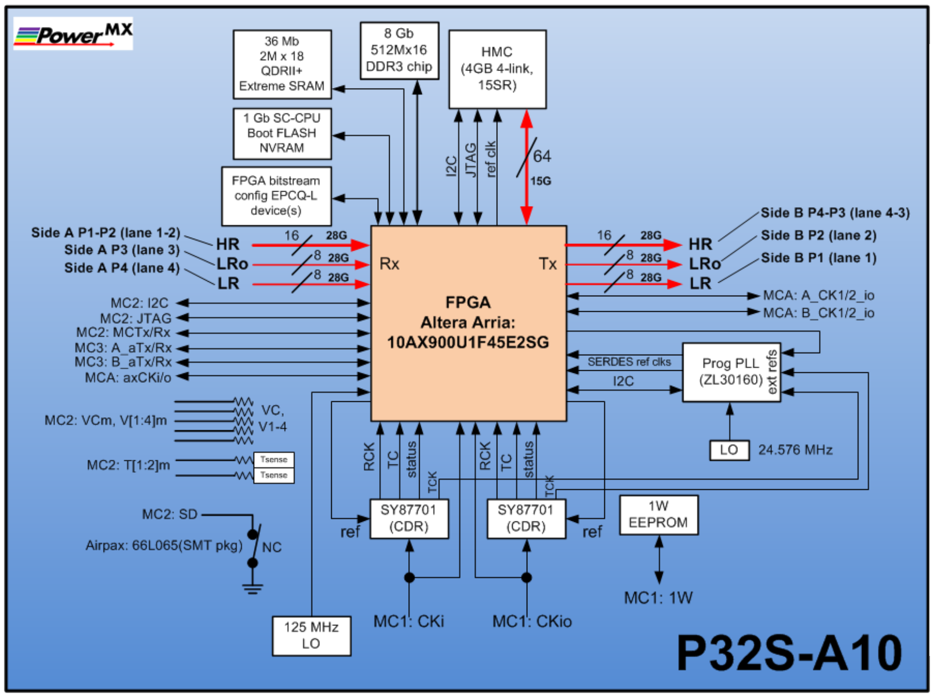

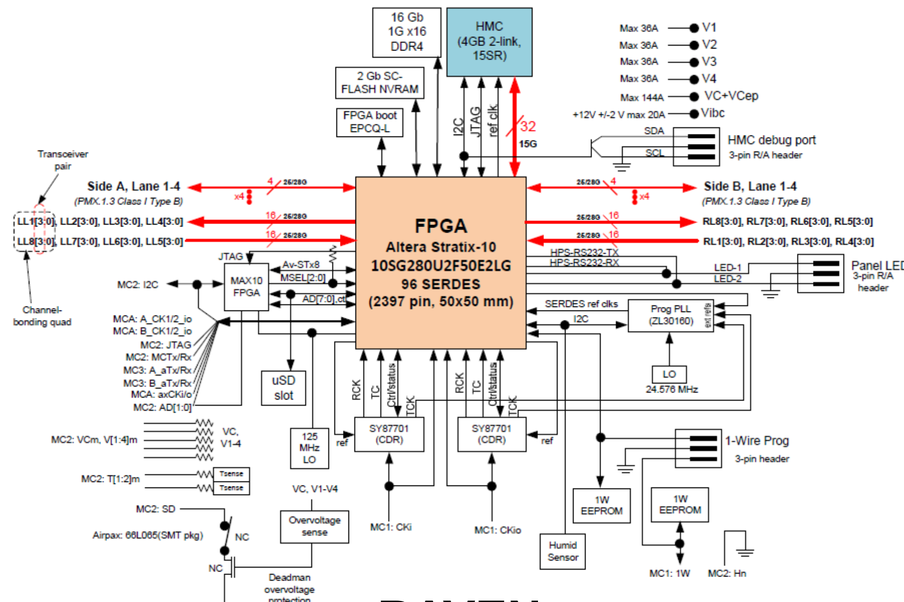

RAVEN

SKA1-Mid.CBF

- Thermal density, total power, and space demands liquid cooling
- NRC is working with CoolIT Systems (Calgary, Canada)
  - Customized thermal plate design
  - Rack manifolds, liquid-liquid or liquid-air heat exchangers
  - Looking at non-water facility liquid delivery using HFE7000 fluid
  - High reliability, redundancy, robustness
  - Dry quick connect/disconnect
  - Leak detection
  - Ethernet monitor and control

Thank you

NRC Correlator Program inquiries:
David.Loop@nrc.ca