NRC·CNRC

Correlator and Beamformer technology for the SKA CSP

Brent Carlson

ngVLA Workshop @ Caltech, April 9, 2015

NRC-Herzberg Astronomy Technology Program

National Research Council Canada | Conseil national de recherches Canada

Canada

Outline

- Overview of re-baselined SKA telescope(s).
- SKA CSP technologies/approaches.
  - Uniboard-2 (JIVE).
  - Redback-x (CSIRO).
  - SKARAB (SKA-SA).
- PowerMX.
- ngVLA correlator possibilities using PowerMX.

Overview of re-baselined telescopes

- SKA1-Mid:
  - 133, 15 m SKA1 dishes + 64 12 m MeerKAT dishes.
  - Band 2 (810 MHz/p), Band 5 (2 x 2.5 GHz/p), Band 1 (700 MHz/p).
  - 64,000 channels pp plus (16X) zoom mode.
  - Central beamforming for pulsar searching (~1000-1500 PSS beams, 300 MHz/beam, 4096 channels) and timing (16 PST beams, 2.5 GHz/beam, 10 MHz channels)
  - Real-time pulsar searching on every PSS beam.
  - Real-time pulsar timing on every PST beam.

SKA CSP technologies: Uniboard-2

- Uniboard-2 (JIVE + collaborators)
  - 24 transceivers on the front, 48 on the back, 12 connect up, 12 connect down.
  - Main aim is Stratix-10, but first protos built with Arria-10.
  - I/O data rate, 960 Gps(I) and 960(O).
  - 1st protos next 1-2 months.

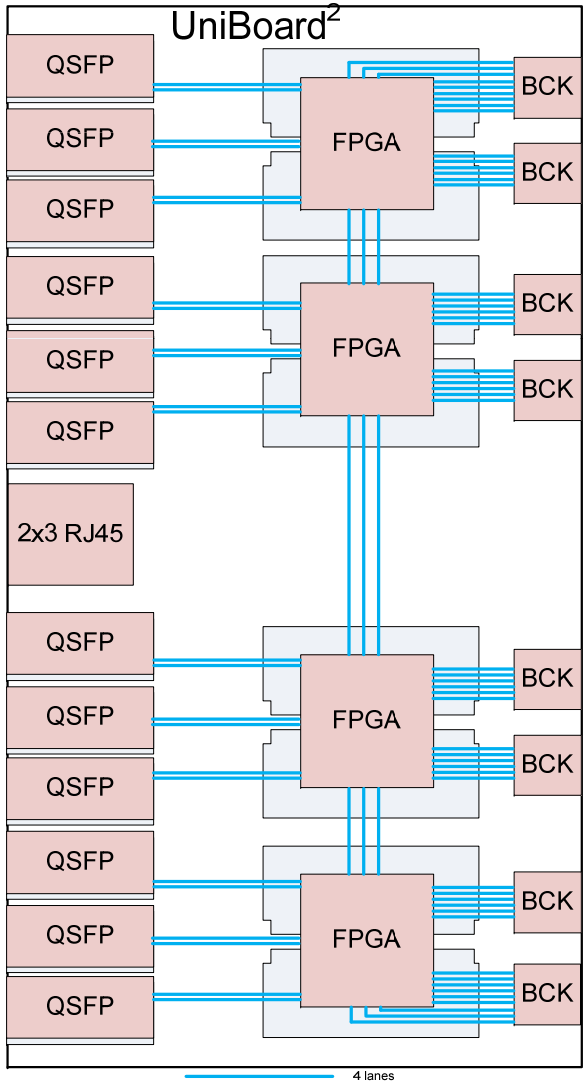

SKA CSP technologies: Uniboard-2

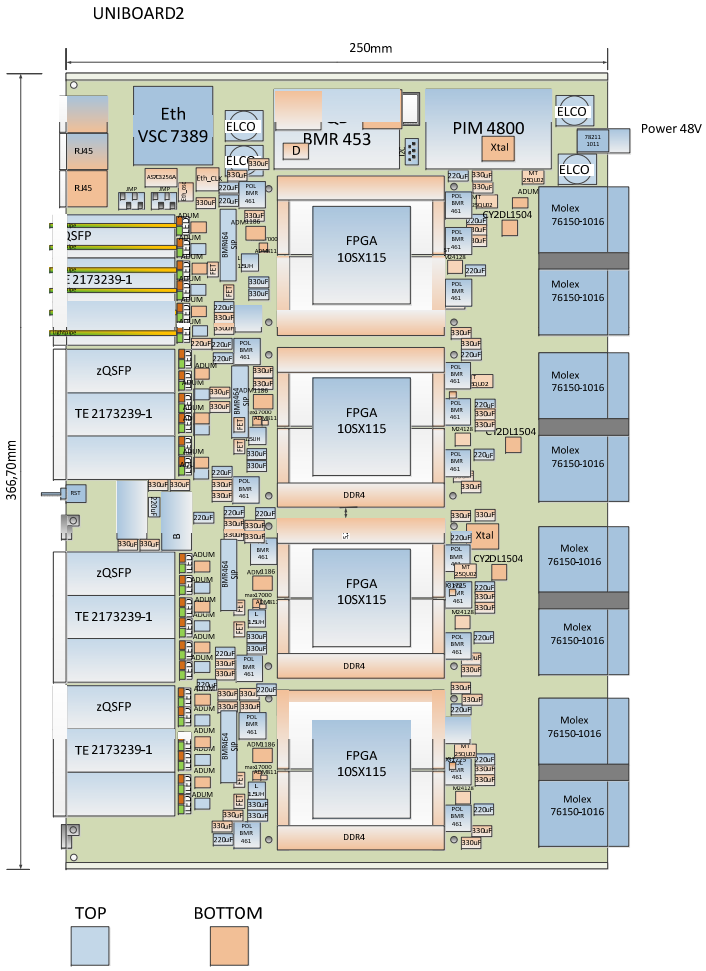

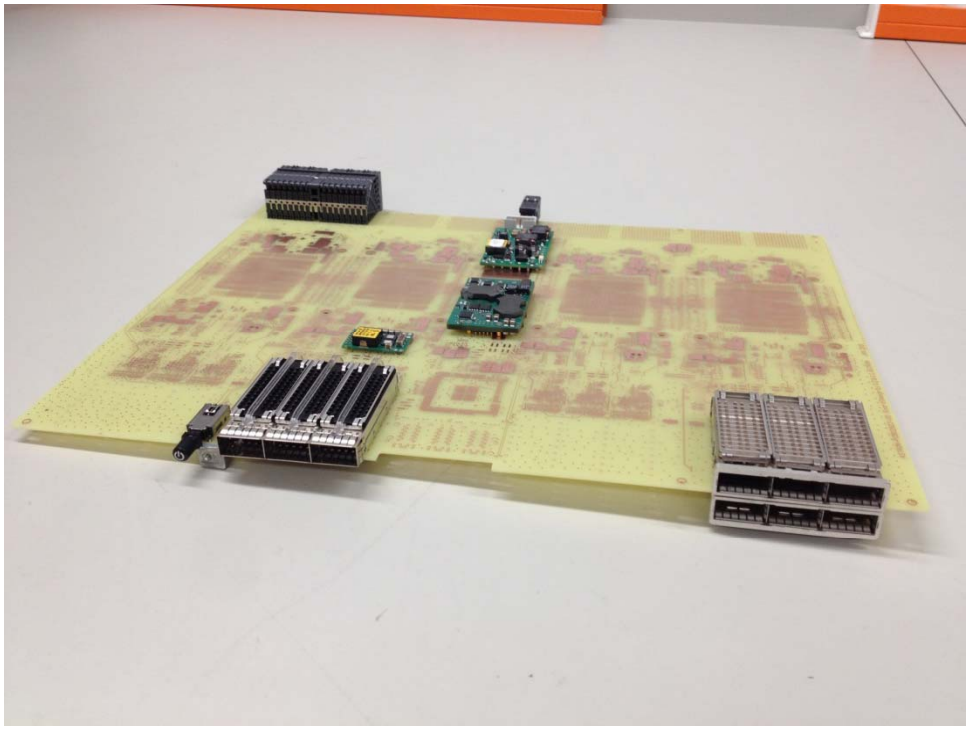

SKA CSP technologies

- Redback-5 (CSIRO), 1U pizza-box (still evolving):

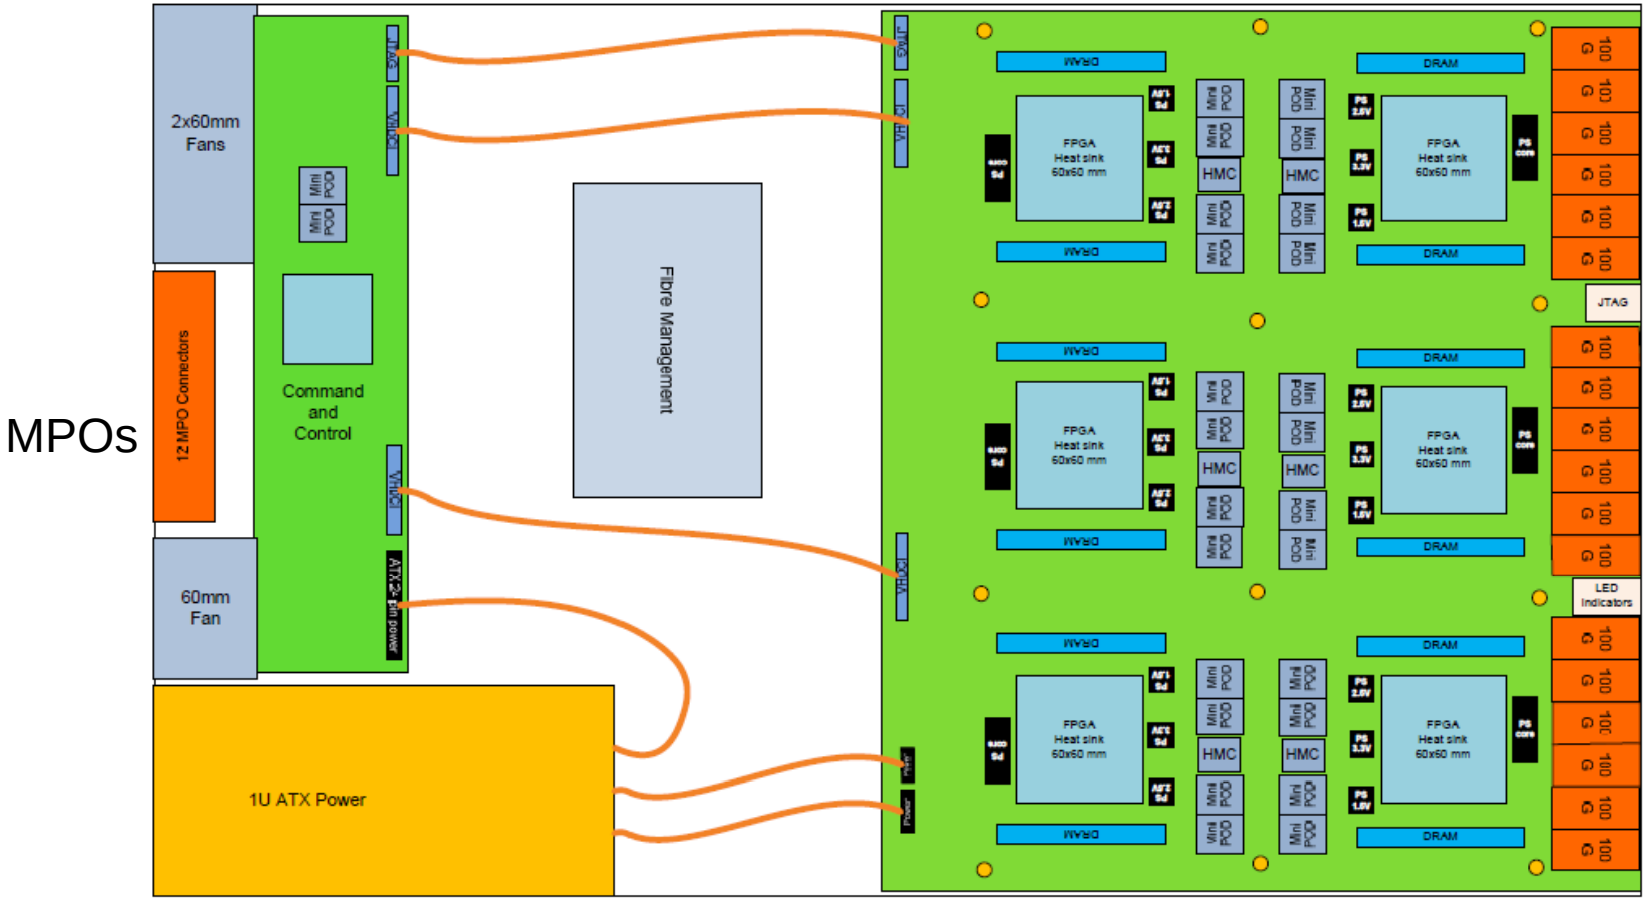


SUNON
SUNON
SUNON
ASKAP Redback

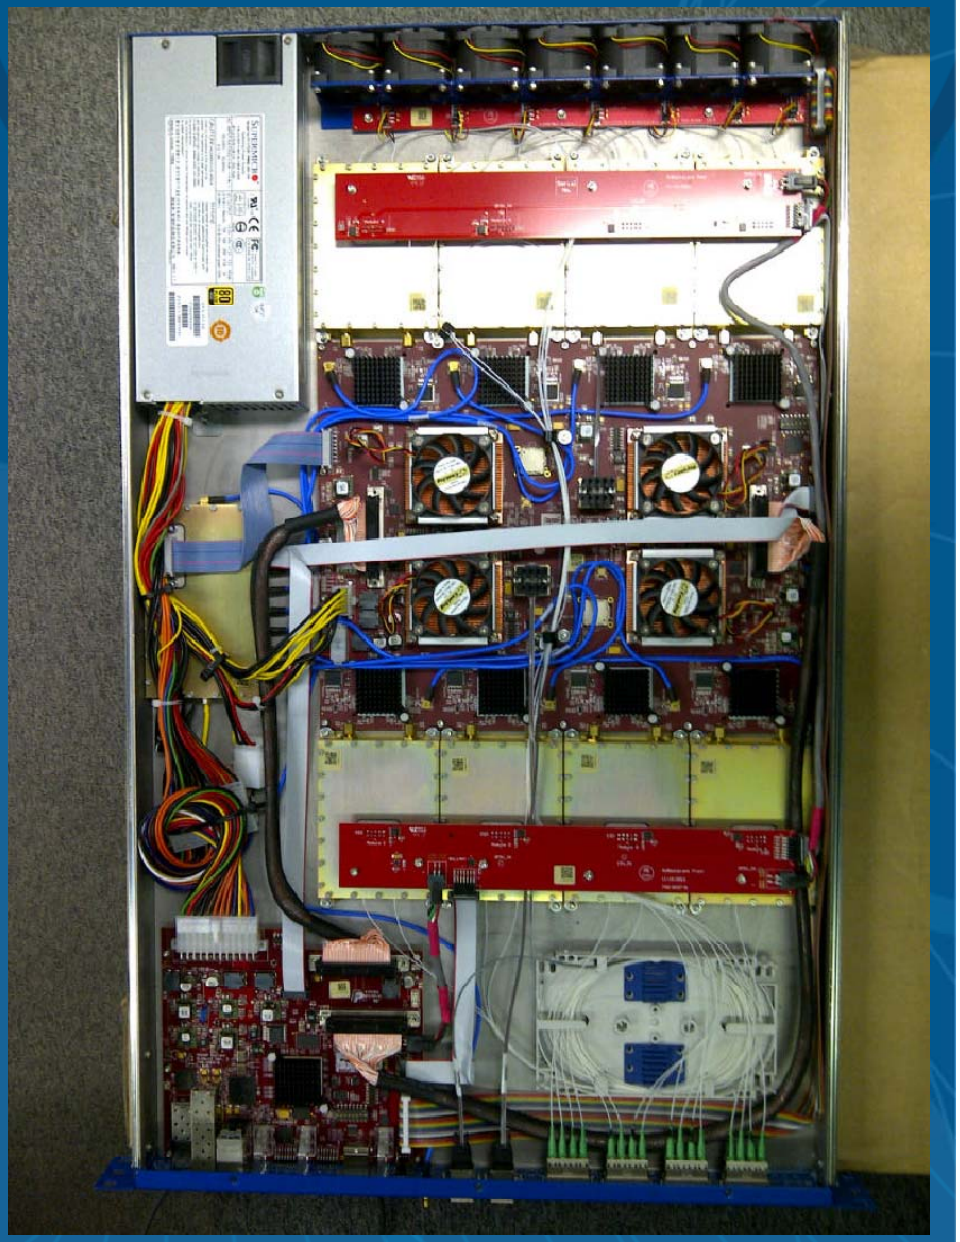

SKA CSP technologies

- SKARAB (SKA-SA)

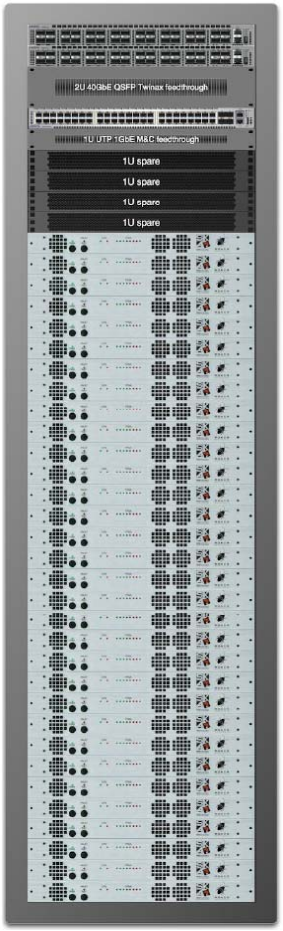

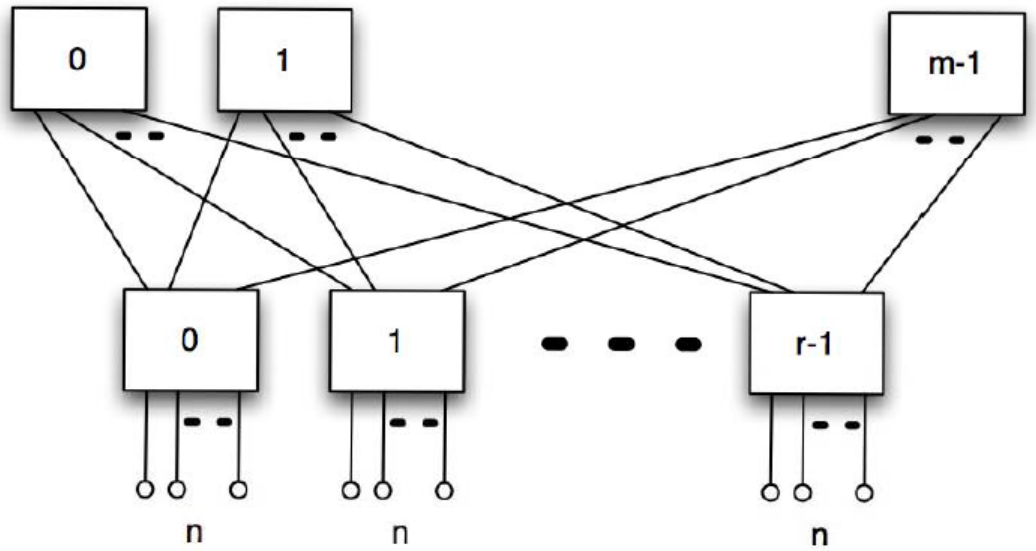

Developed for MeerKAT.
Single Xilinx Virtex-7 FPGA + 40 GbE per pizza box
CASPER tool flow
Scalable distributed switch architecture.

SKA CSP technologies

- PowerMX (NRC and some others)
  - Motherboard/mezzanine approach. Both processing and I/O.
  - “Tons” of I/O and performance for “future-proofness”.
  - Series of open-public specifications to which anyone can build to, with the goal of plug-n-play compatibility.
  - Inspired by CASPER…spurred on by COTS solution competition.
    - Ultimate goal: Buy some PowerMX products from COTS providers.

PowerMX

- Motivation:
  - Future proof…minimize cost/time of technology design cycle, allow for cross-generation mix/match plug-n-play of multi-vendor modules.
  - Compatibility + fuel for innovation.
  - Add new specs as required.
  - Create a framework within which products are developed and plug-n-play compatible.
  - *Facilitate access/connectivity of FPGAs and other processing devices.*

PowerMX PMX.1 Base Specification

- See www.powermx.org open, public, freely available for use.

- Basic concepts:
  - Motherboard (power, M&C) + processing and I/O mezzanines. With coherent clocks to I/O mezzanines…for ADC cards.
  - Primarily SERDES I/O.
  - Use 4 mm FCI Meg-Array mezzanine connectors.
  - Full 4-site motherboard: 384 serial I/Os (@28G ea ~= 10.75 Tbps).
  - Spec allows for “hardened”.

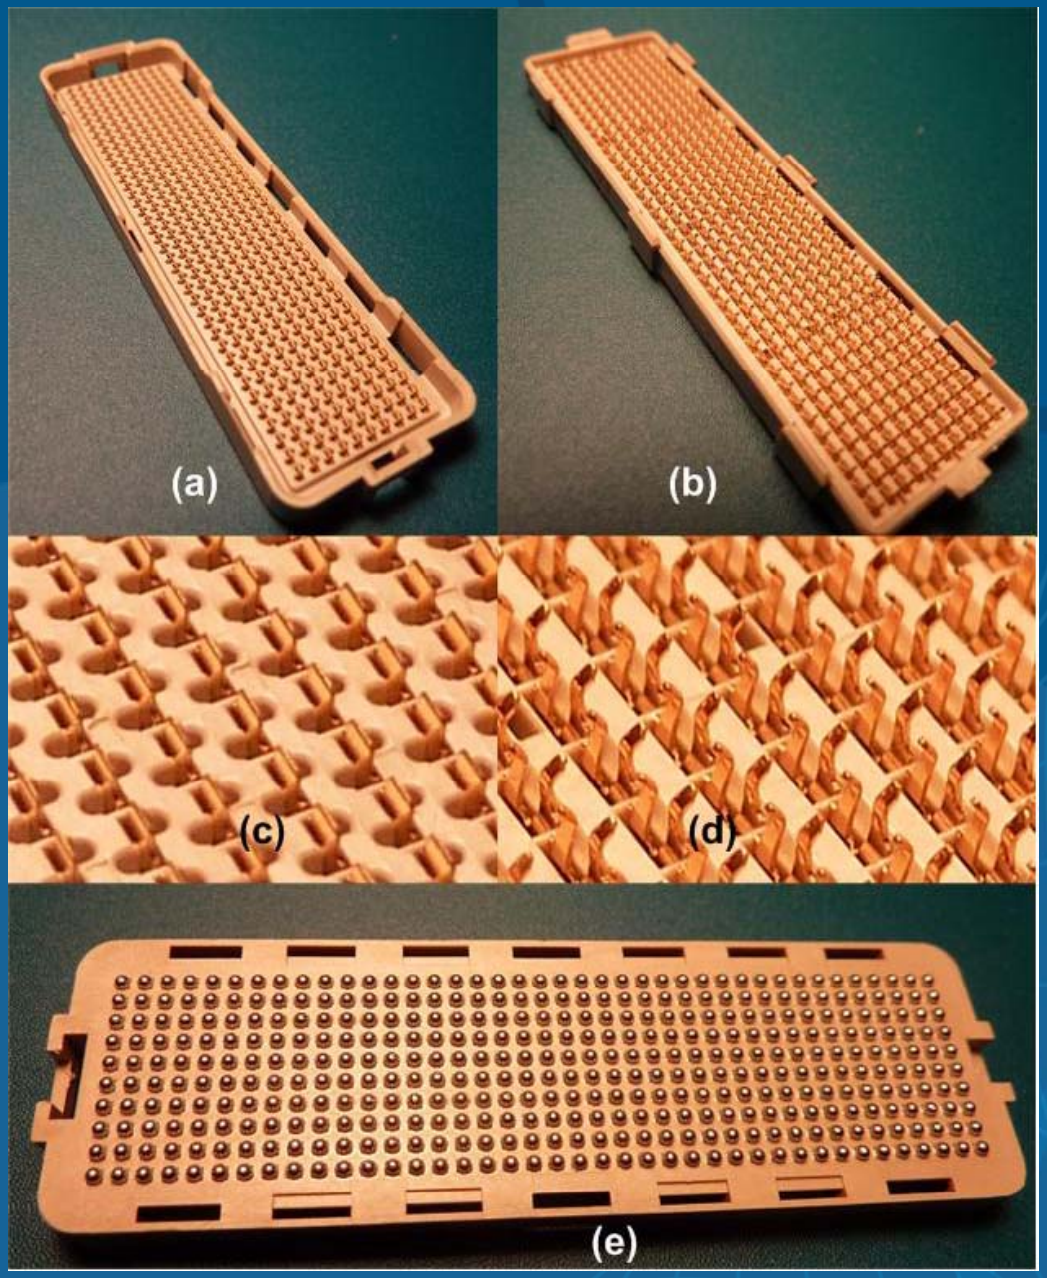
(a)
(b)
(c)
(d)
(e)

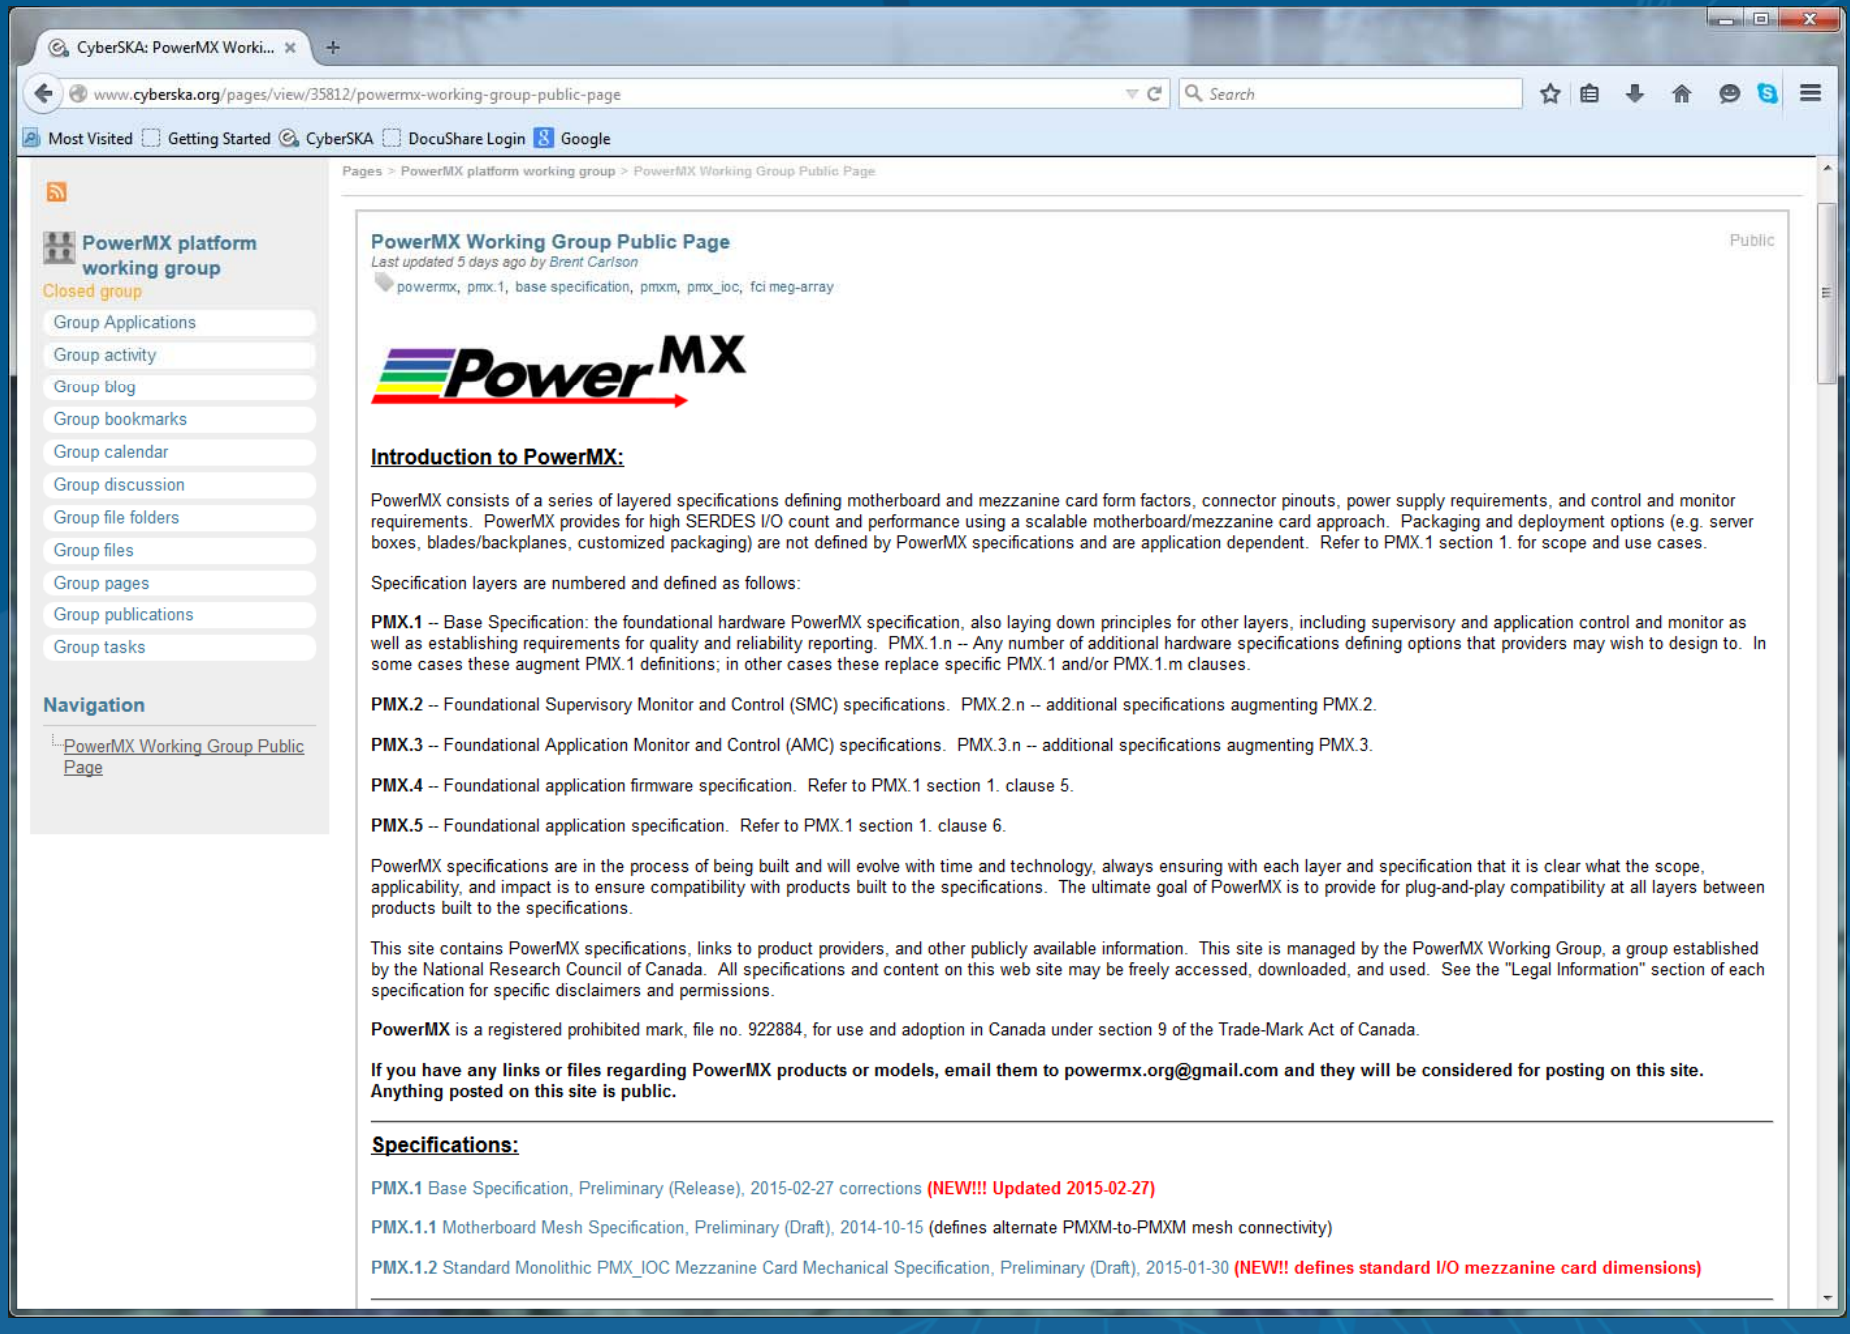
CyberSKA: PowerMX Worki...
www.cyberska.org/pages/view/35812/powermx-working-group-public-page
Search
Most Visited
Getting Started
CyberSKA
DocuShare Login
Google
Pages > PowerMX platform working group > PowerMX Working Group Public Page
PowerMX platform working group
Closed group
Group Applications
Group activity
Group blog
Group bookmarks
Group calendar
Group discussion
Group file folders
Group files
Group pages
Group publications
Group tasks
Navigation
PowerMX Working Group Public Page
PowerMX Working Group Public Page
Last updated 5 days ago by Brent Carlson
powermx, pmx.1, base specification, pmxm, pmx_ioc, fci meg-array
Public
PowerMX
Introduction to PowerMX:
PowerMX consists of a series of layered specifications defining motherboard and mezzanine card form factors, connector pinouts, power supply requirements, and control and monitor requirements. PowerMX provides for high SERDES I/O count and performance using a scalable motherboard/mezzanine card approach. Packaging and deployment options (e.g. server boxes, blades/backplanes, customized packaging) are not defined by PowerMX specifications and are application dependent. Refer to PMX.1 section 1. for scope and use cases.
Specification layers are numbered and defined as follows:
PMX.1 -- Base Specification: the foundational hardware PowerMX specification, also laying down principles for other layers, including supervisory and application control and monitor as well as establishing requirements for quality and reliability reporting. PMX.1.n -- Any number of additional hardware specifications defining options that providers may wish to design to. In some cases these augment PMX.1 definitions; in other cases these replace specific PMX.1 and/or PMX.1.m clauses.
PMX.2 -- Foundational Supervisory Monitor and Control (SMC) specifications. PMX.2.n -- additional specifications augmenting PMX.2.
PMX.3 -- Foundational Application Monitor and Control (AMC) specifications. PMX.3.n -- additional specifications augmenting PMX.3.
PMX.4 -- Foundational application firmware specification. Refer to PMX.1 section 1. clause 5.
PMX.5 -- Foundational application specification. Refer to PMX.1 section 1. clause 6.
PowerMX specifications are in the process of being built and will evolve with time and technology, always ensuring with each layer and specification that it is clear what the scope, applicability, and impact is to ensure compatibility with products built to the specifications. The ultimate goal of PowerMX is to provide for plug-and-play compatibility at all layers between products built to the specifications.
This site contains PowerMX specifications, links to product providers, and other publicly available information. This site is managed by the PowerMX Working Group, a group established by the National Research Council of Canada. All specifications and content on this web site may be freely accessed, downloaded, and used. See the "Legal Information" section of each specification for specific disclaimers and permissions.
PowerMX is a registered prohibited mark, file no. 922884, for use and adoption in Canada under section 9 of the Trade-Mark Act of Canada.
If you have any links or files regarding PowerMX products or models, email them to powermx.org@gmail.com and they will be considered for posting on this site. Anything posted on this site is public.
Specifications:
PMX.1 Base Specification, Preliminary (Release), 2015-02-27 corrections (NEW!!! Updated 2015-02-27)
PMX.1.1 Motherboard Mesh Specification, Preliminary (Draft), 2014-10-15 (defines alternate PMXM-to-PMXM mesh connectivity)
PMX.1.2 Standard Monolithic PMX_IOC Mezzanine Card Mechanical Specification, Preliminary (Draft), 2015-01-30 (NEW!! defines standard I/O mezzanine card dimensions)

PowerMX PMX.1 Base Specification

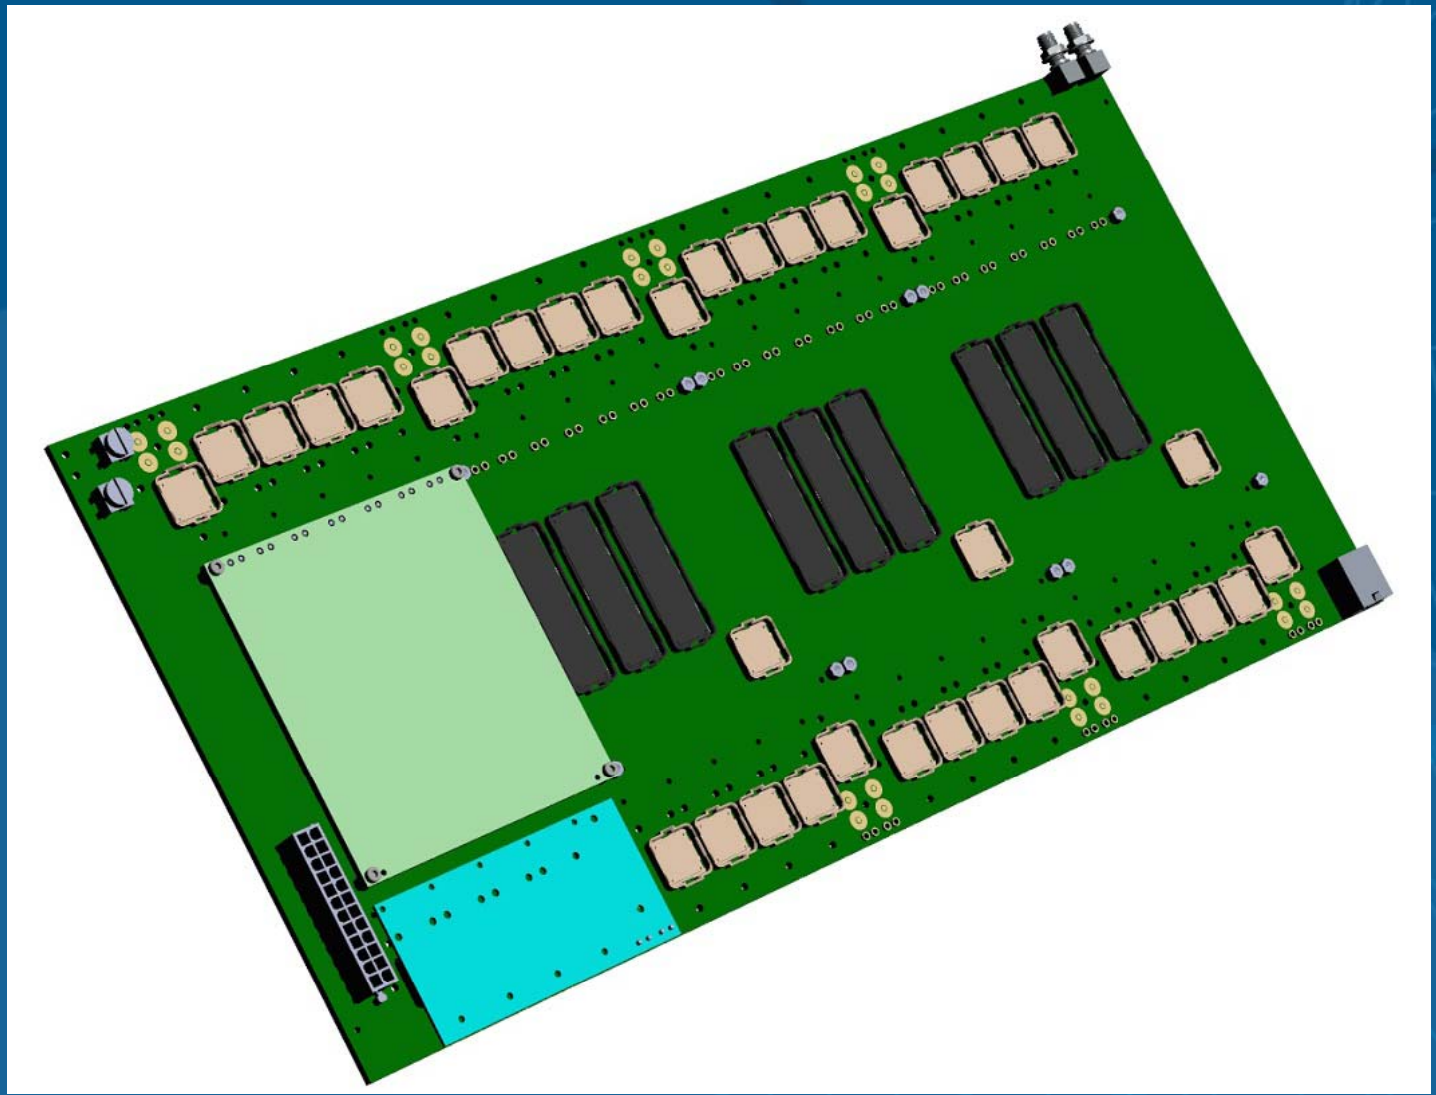

PowerMX PMX.1 Base Specification

PowerMX

PMX.1

Base Specification

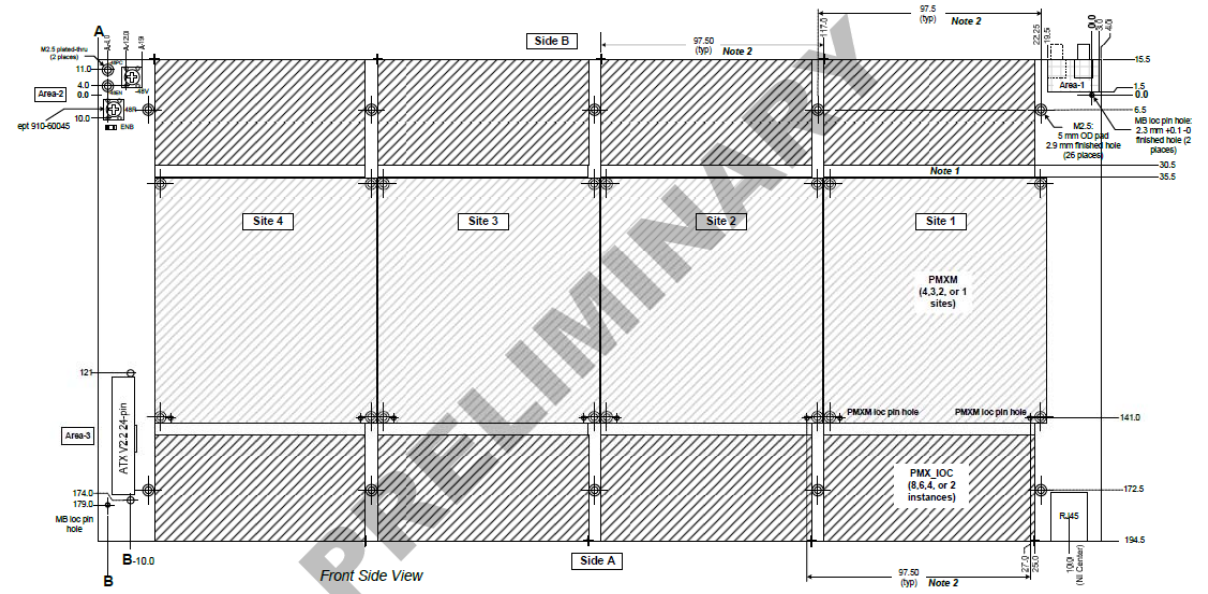

Figure 5-1 PowerMX motherboard dimensions. Coordinates A and B depend on the number of sites provided and are defined in Table 5-1.

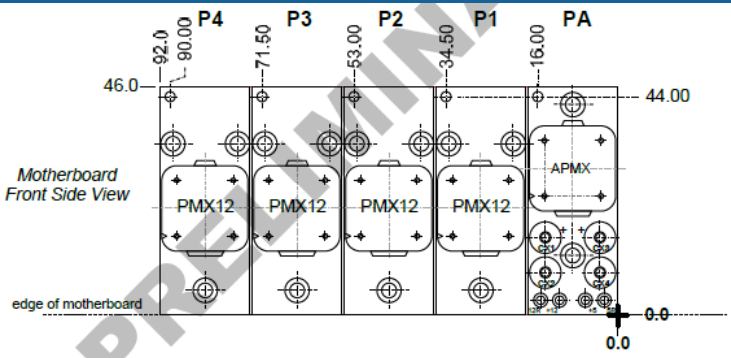

Figure 5-3 PMX_IOC footprint on motherboard showing key coordinates and nomenclature. Note that APMX and PMX12 footprint dimensions are 18.0 mm x 46.0 mm with an 18.50 mm intra PMX_IOC module-to-module pitch.

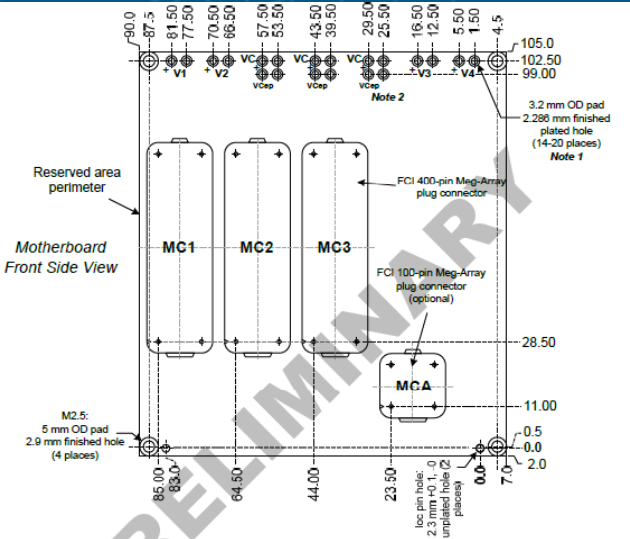

Figure 5-6 PMXM footprint on the motherboard. The "reserved area perimeter" indicates the perimeter within which only those components shown in this drawing may be located on the Front Side of the motherboard.

NRC·CNRC

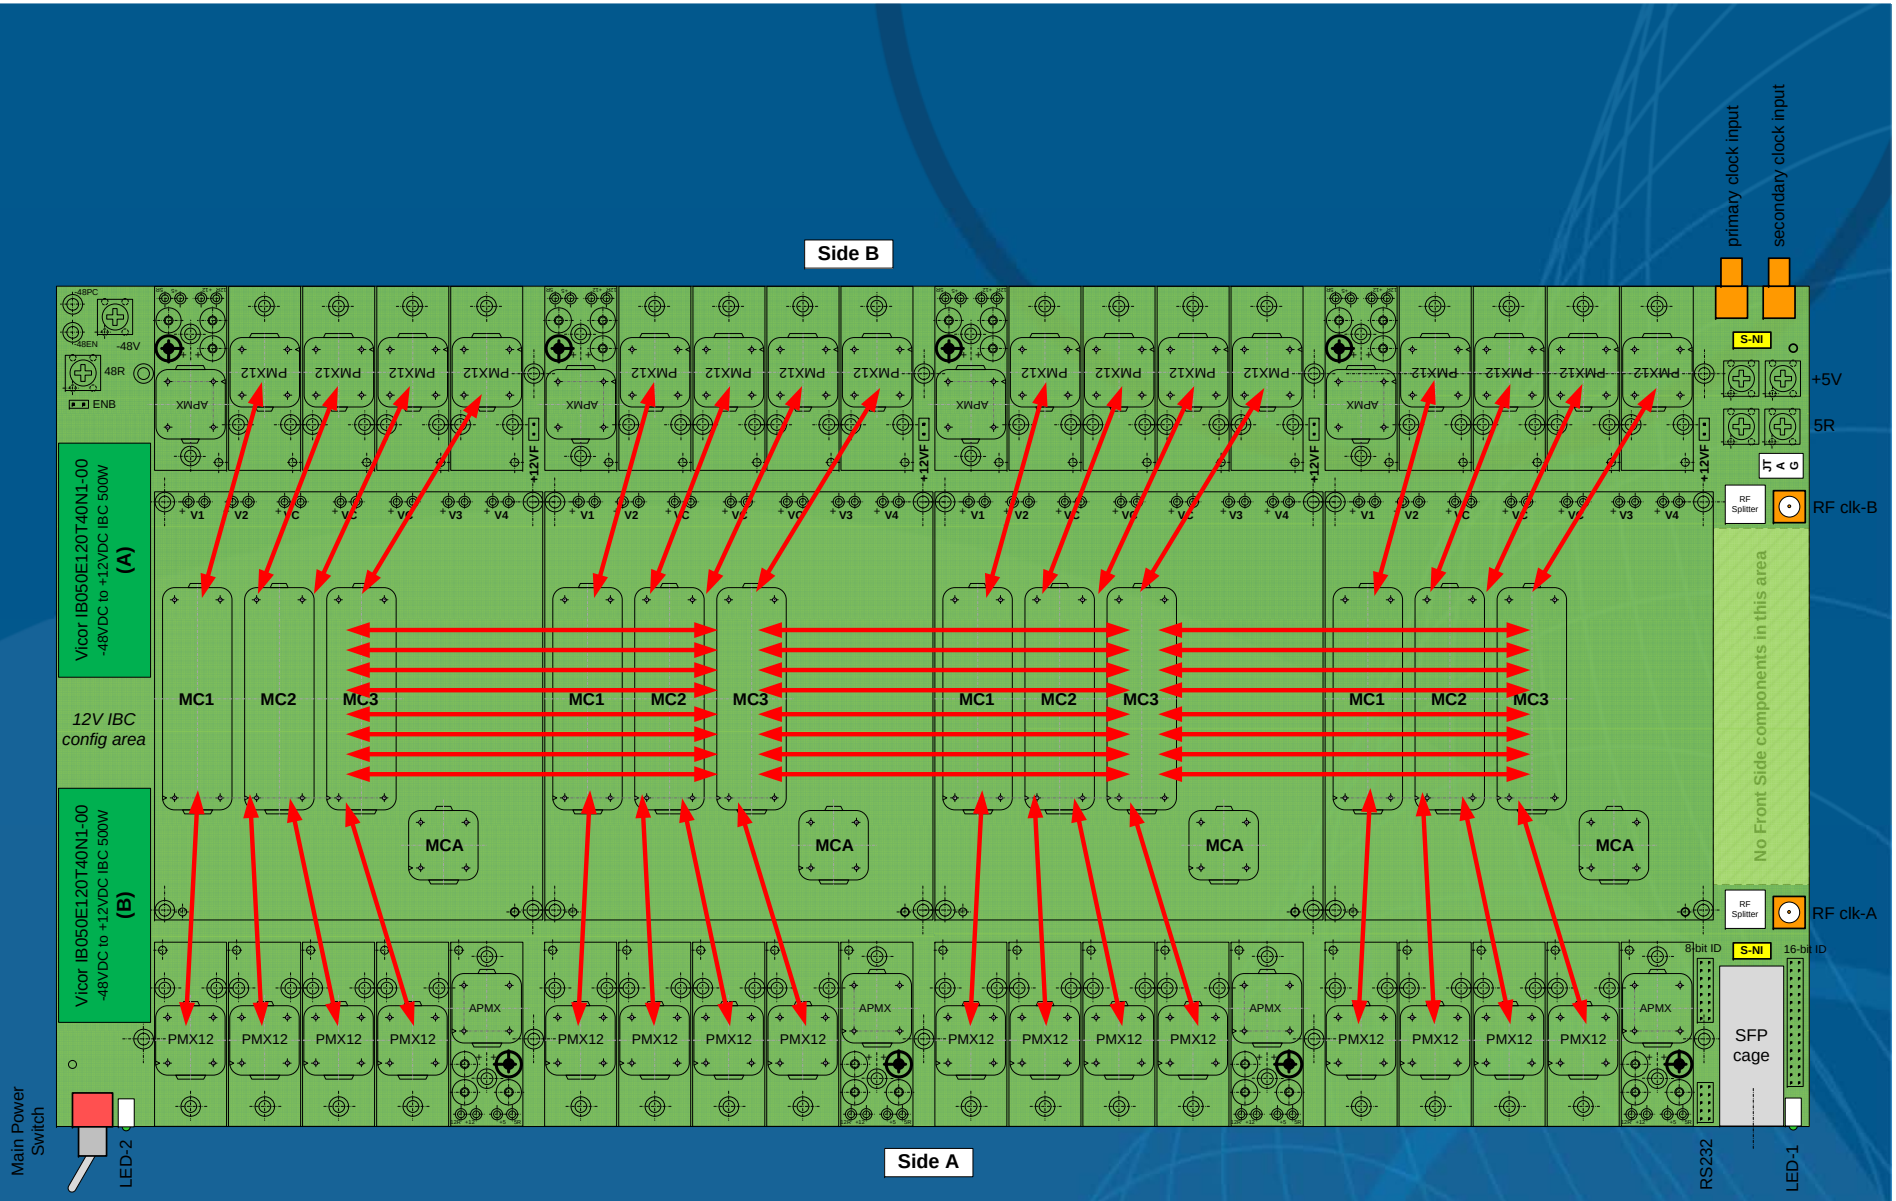
Side B
primary clock input
secondary clock input
S-NI
+5V
5R
RF clk-B
RF clk-A
Vicor IB050E120T40N1-00
-48VDC to +12VDC IBC 500W
(A)
Vicor IB050E120T40N1-00
-48VDC to +12VDC IBC 500W
(B)
12V IBC
config area
MC1
MC2
MC3
MCA
PMX12
APMX
+12VF
V1
V2
VC
V3
V4
RF Splitter
No Front Side components in this area
8-bit ID
16-bit ID
SFP
cage
RS232
LED-1
LED-2
Main Power
Switch
Side A

PowerMX PMX.1.1 Mesh Specification

- Alternate motherboard PMXM-to-PMXM mesh connectivity.

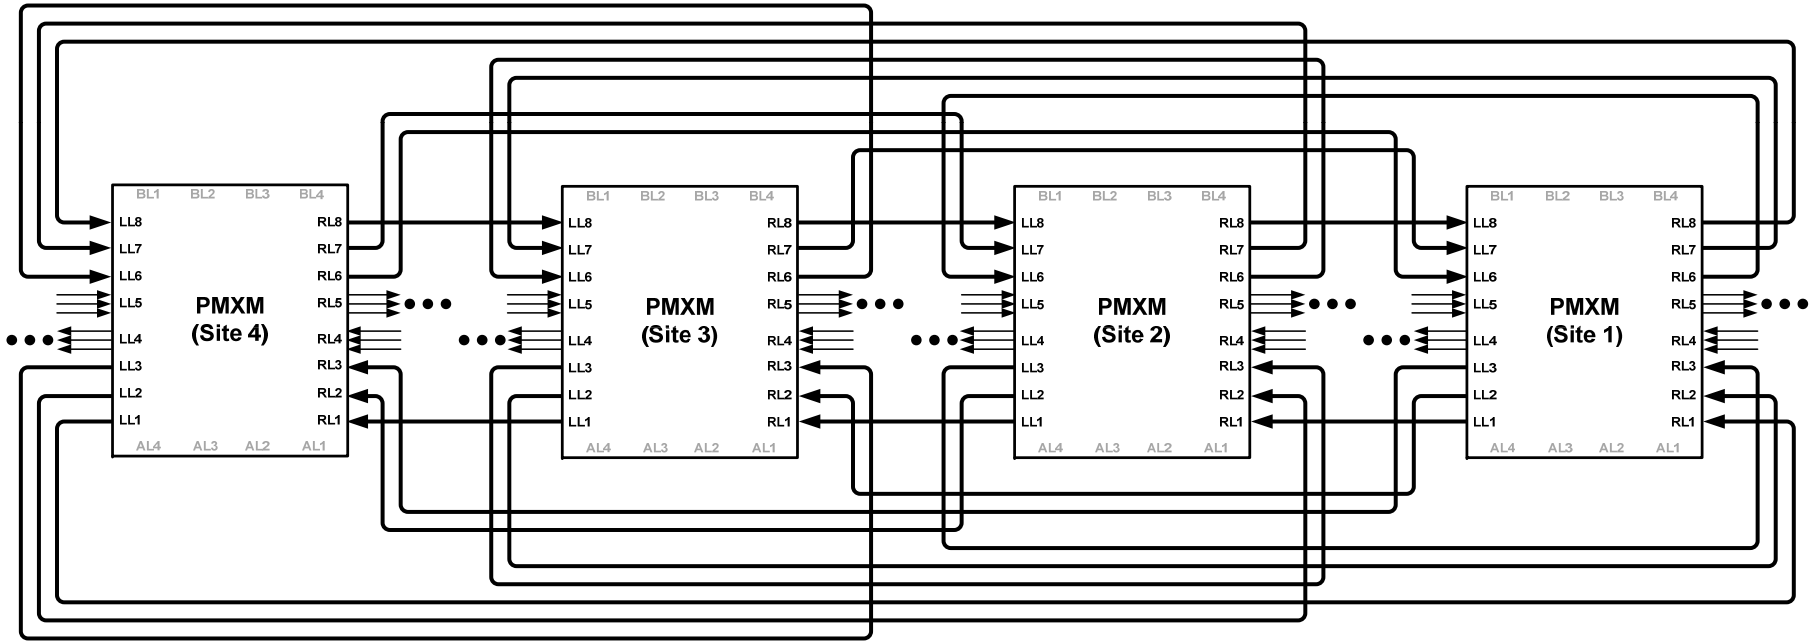

PowerMX PMX.1 Base Specification

- +Establishes base level of quality and reliability reporting.
- Other layers in other documents:
  - Other H/W layers (PMX.1.x)
  - SMC: Supervisory M&C (PMX.2.x)
  - AMC : Application M&C (PMX.3.x)
  - PMX.4.x: Application F/W
  - PMX.5.x: Applications…

PowerMX SMC and AMC

- **SMC**="supervisory monitor and control".
  - PowerMX logical addressing…with implementation-specific extensions.
  - GigE UDP/IP packets for comms to motherboard "SMC controller"…provides access to all I/O non-GigE comm paths and V+T+P control/monitor.
  - PMX.2—simple layer (all implement). PMX.2.n—more soph. layers.

- **AMC**="application monitor and control".
  - Memory-mapped (poke/peek) access to "applications" on mezzanine cards.
  - GigE UDP/IP packets.
  - PMX.3—simple layer (all implement). PMX.3.n—more soph. layers.

- Any processor on the network, via any Ethernet path, can be an AMC or SMC host…even a mezzanine card. <u>All access SMC and AMC servers in a standard way.</u>

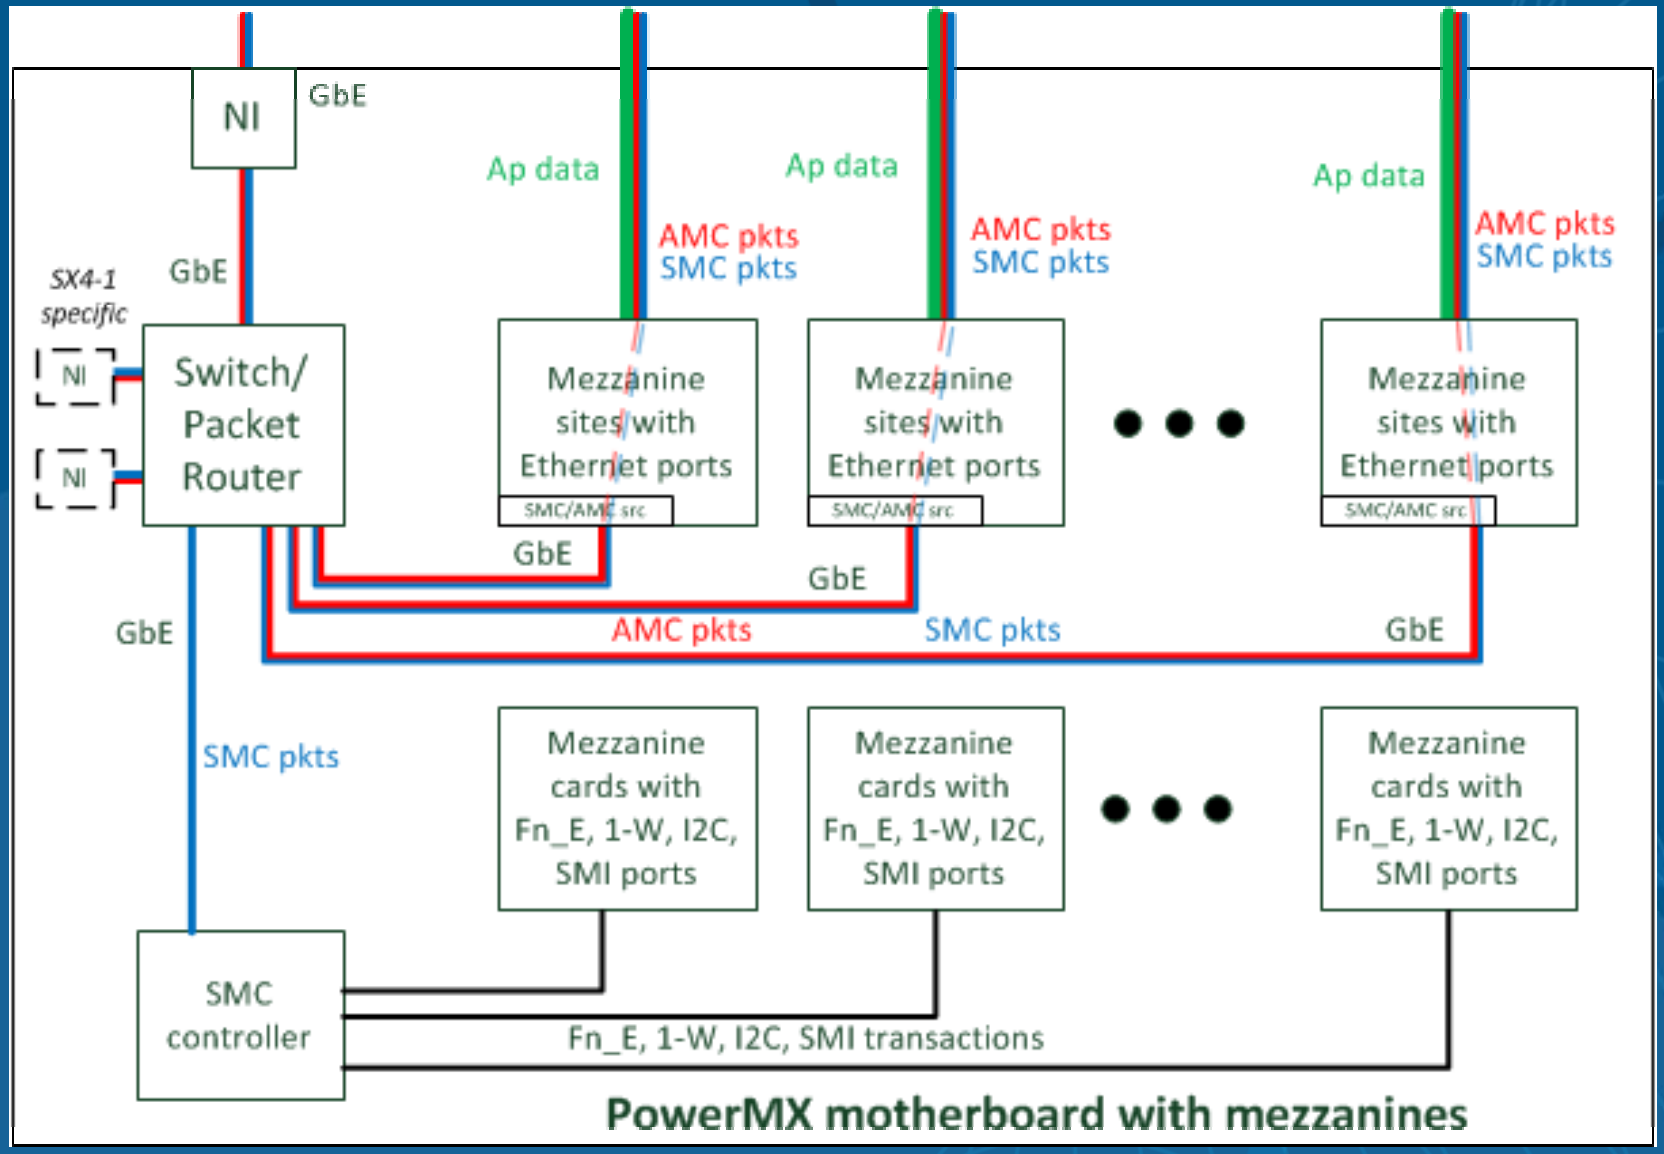
NI
GbE
Ap data
AMC pkts
SMC pkts
Ap data
AMC pkts
SMC pkts
Ap data
AMC pkts
SMC pkts
SX4-1
specific
GbE
NI
NI
Switch/
Packet
Router
Mezzanine
sites with
Ethernet ports
SMC/AMC src
Mezzanine
sites with
Ethernet ports
SMC/AMC src
Mezzanine
sites with
Ethernet ports
SMC/AMC src
GbE
GbE
GbE
AMC pkts
SMC pkts
GbE
SMC pkts
Mezzanine
cards with
Fn_E, 1-W, I2C,
SMI ports
Mezzanine
cards with
Fn_E, 1-W, I2C,
SMI ports
Mezzanine
cards with
Fn_E, 1-W, I2C,
SMI ports
SMC
controller
Fn_E, 1-W, I2C, SMI transactions
PowerMX motherboard with mezzanines

PowerMX PMX.1.2 Standard I/O Module Specification

- Defines standard PMX_IOC (I/O) module dimensions.
- Establishes certainty for compatibility between PCBAs and enclosures.
- Optional…don't need to design to this spec, but useful to do so.

PowerMX PMX.1.2 Specification—standard I/O form factors

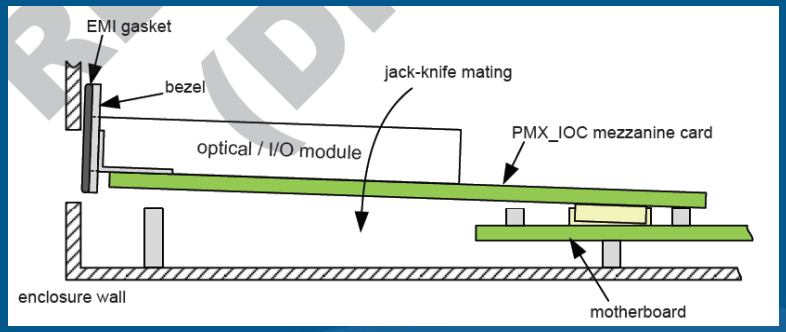

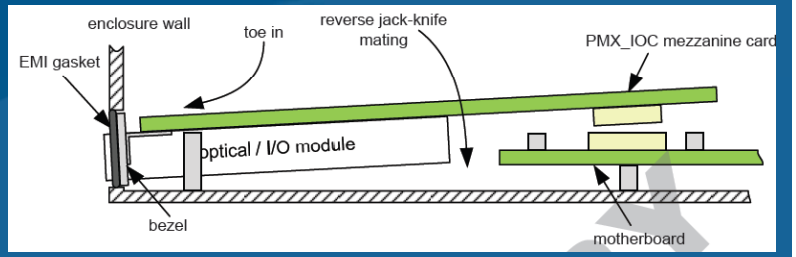

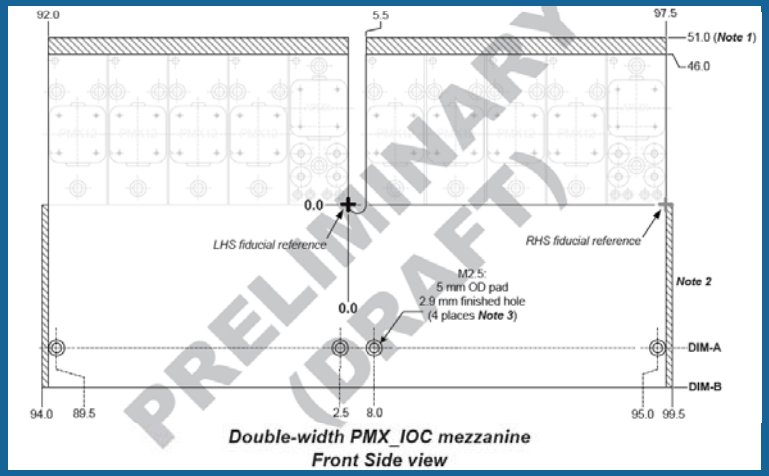

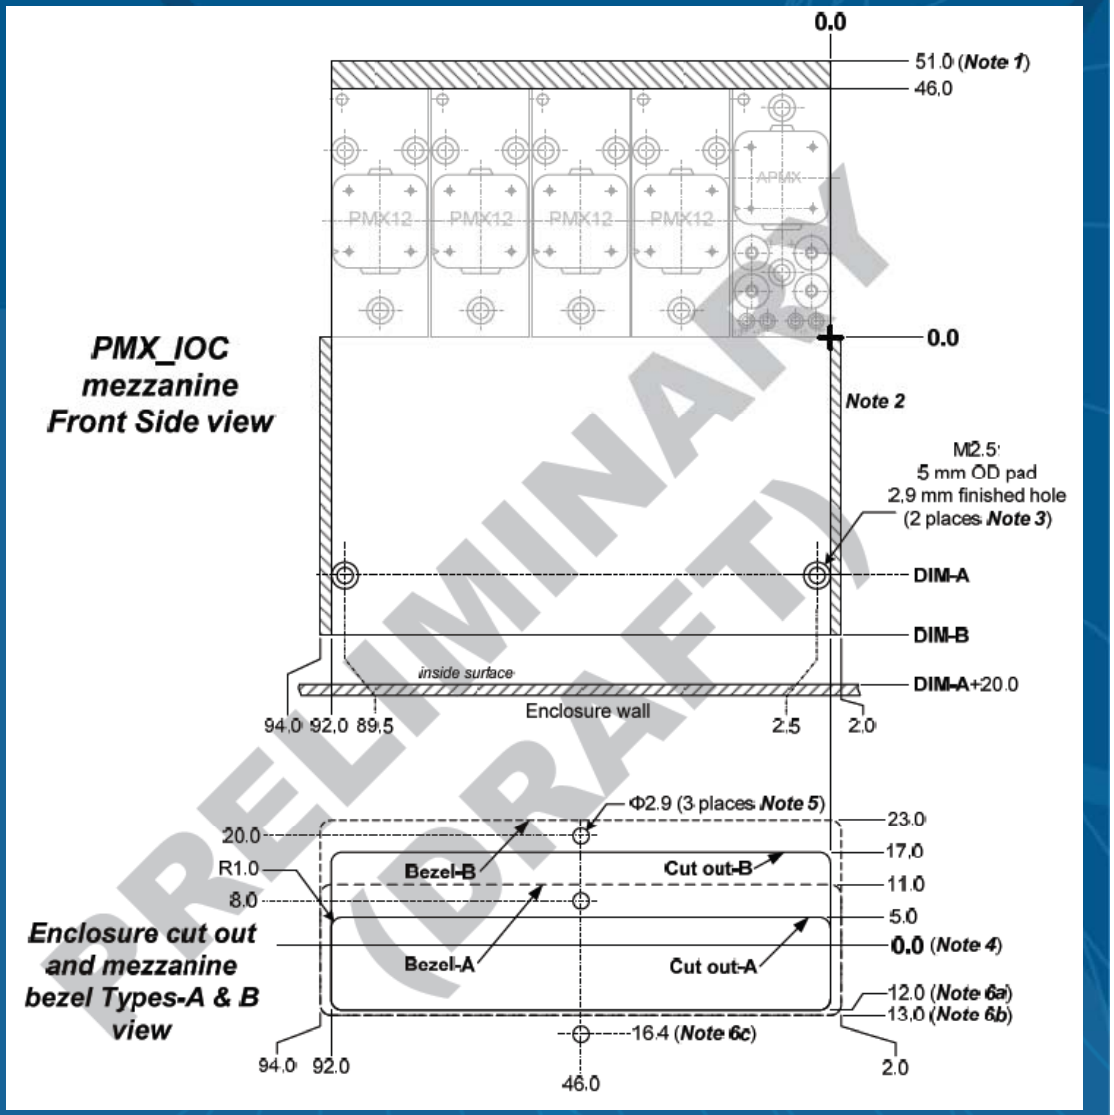

Packaging/Use Examples: Collage of possibilities

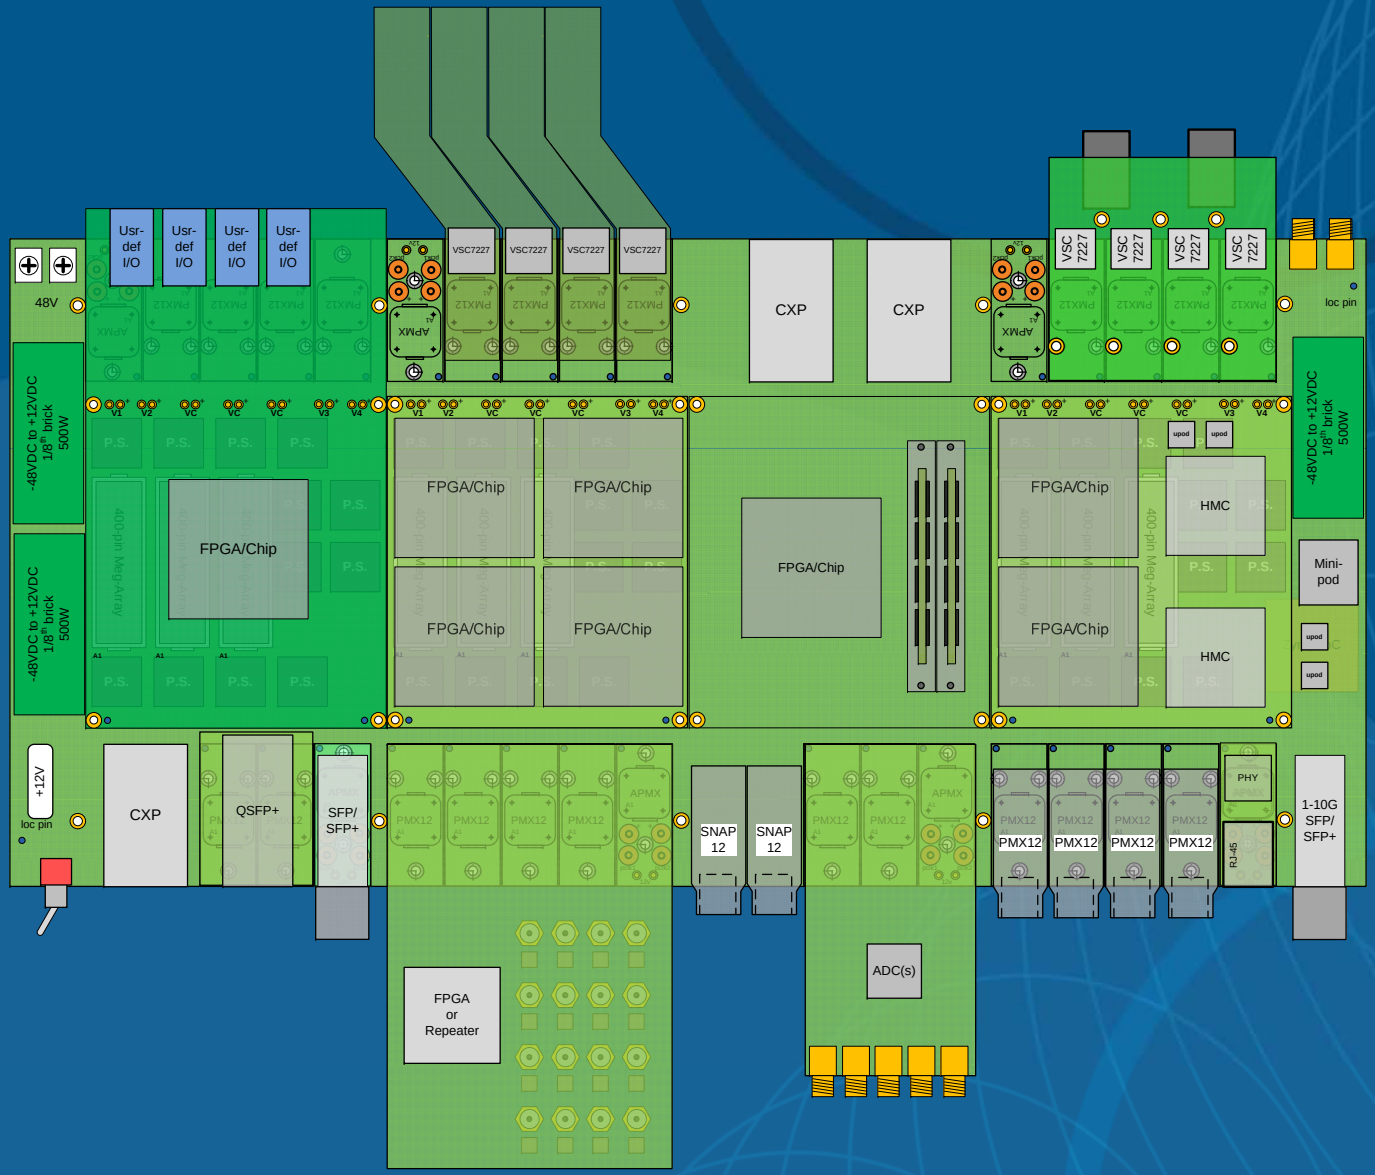

Packaging/Use Example, Band 3 (1.5-4 GHz) cryo-PAF DBE/Beamformer: 96 RF inputs, UofC 4-bit 10G low-power ADC, TDF/TDBF ASICs, 500/1000 MHz beamformed BW, 18 dual-pol beams, double-shielded box, ~650 W.

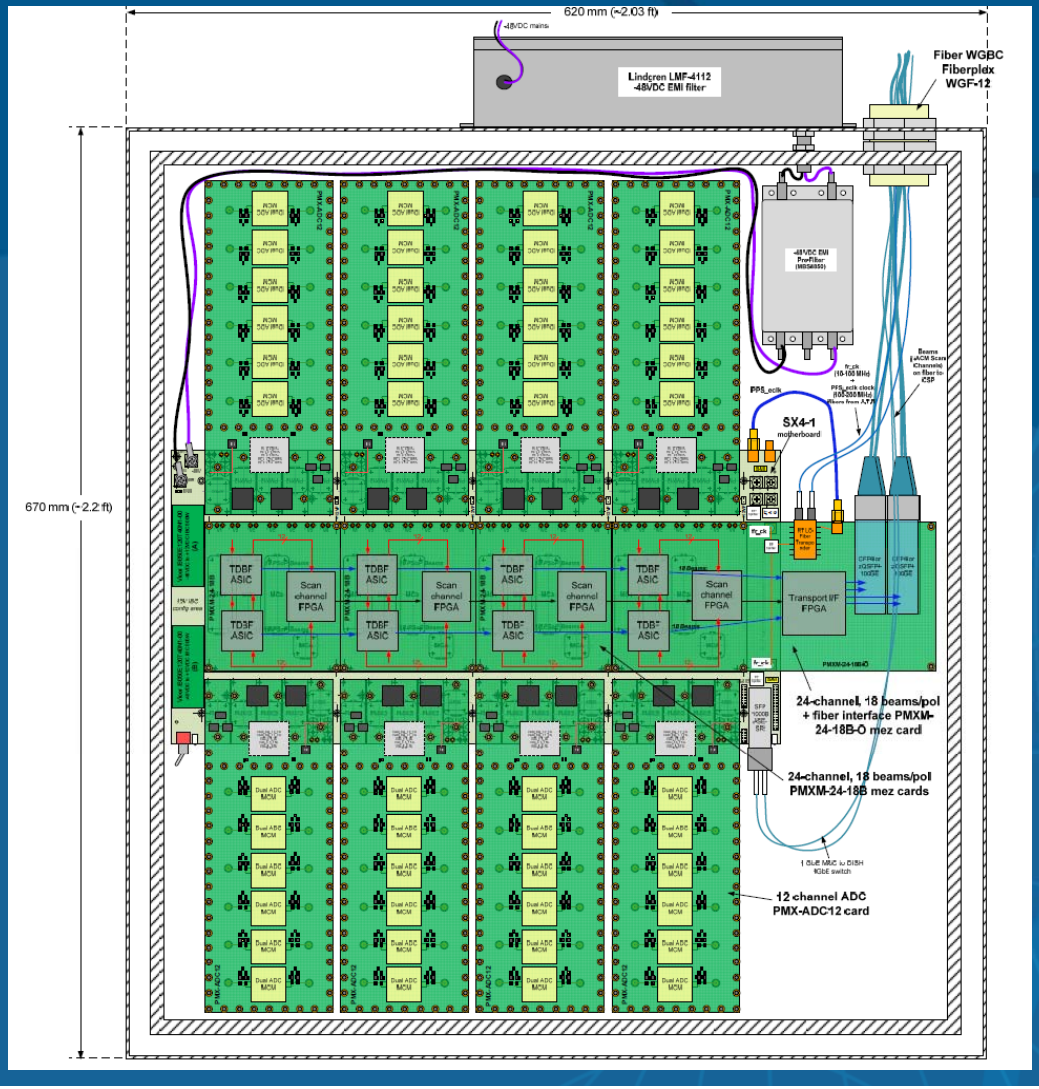

Packaging/Use Examples: Pizza box

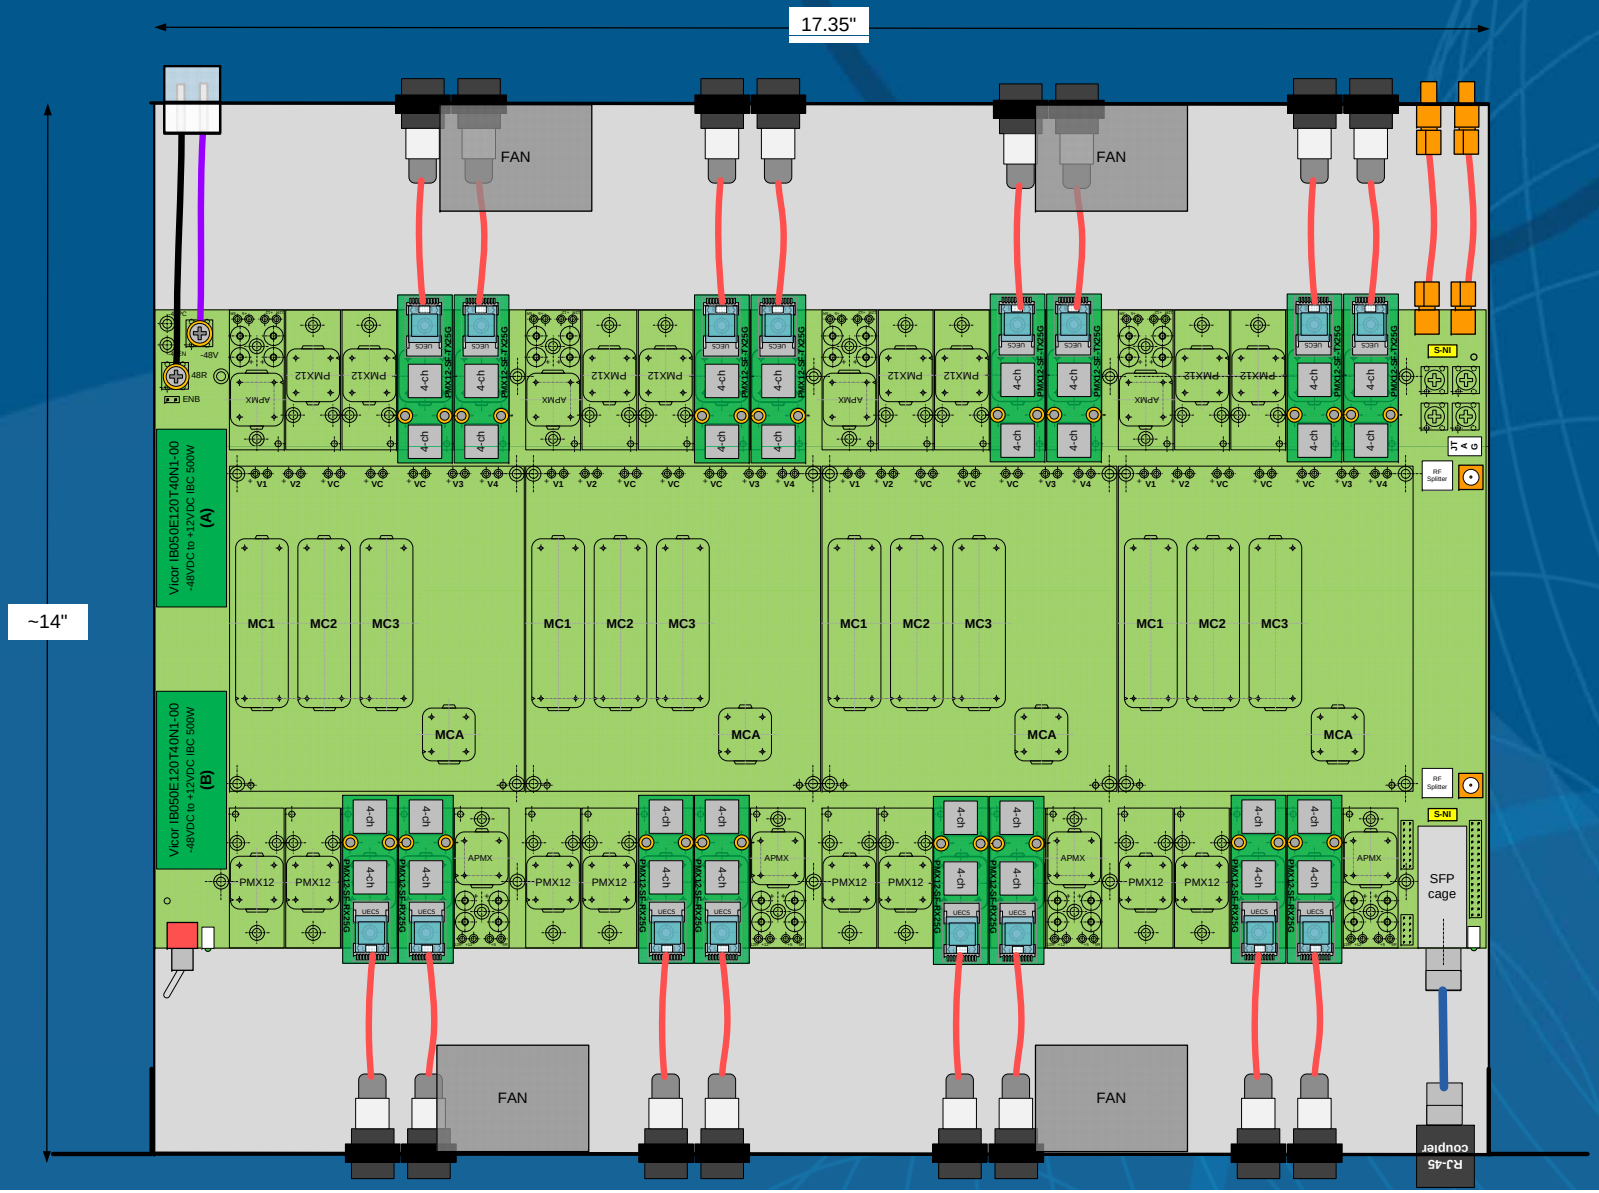

Packaging/Use Examples: Multi-board blade

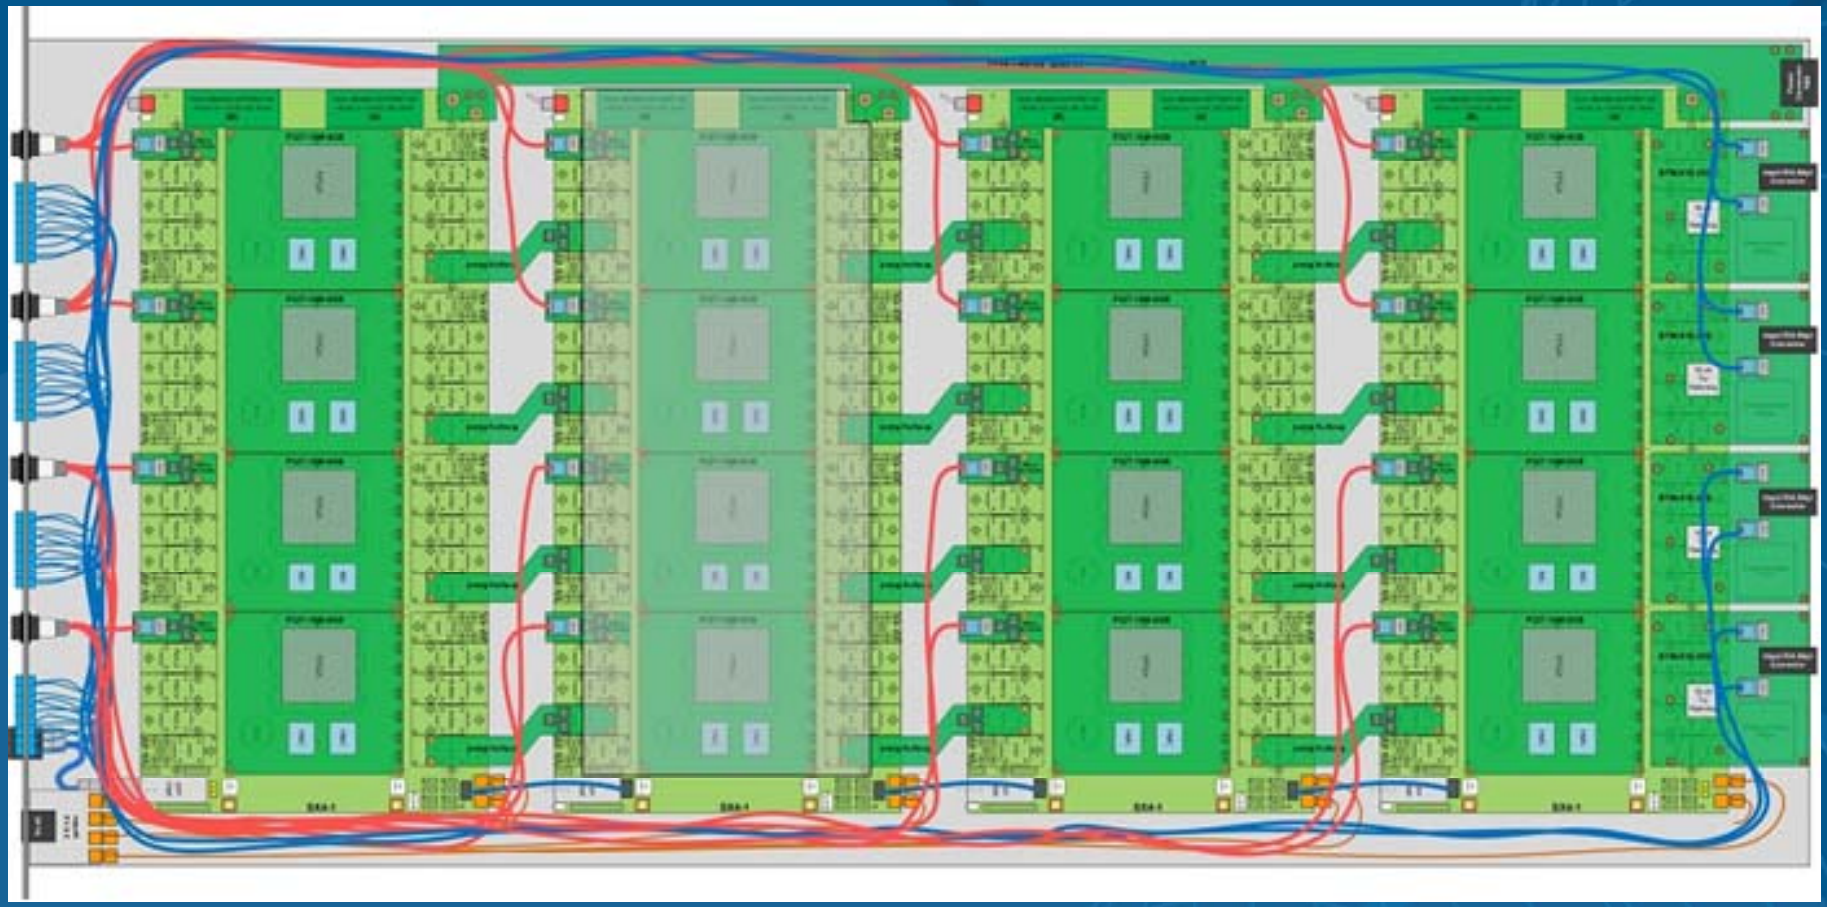

NRC's "SX4-1" motherboard: ready for proto fab.

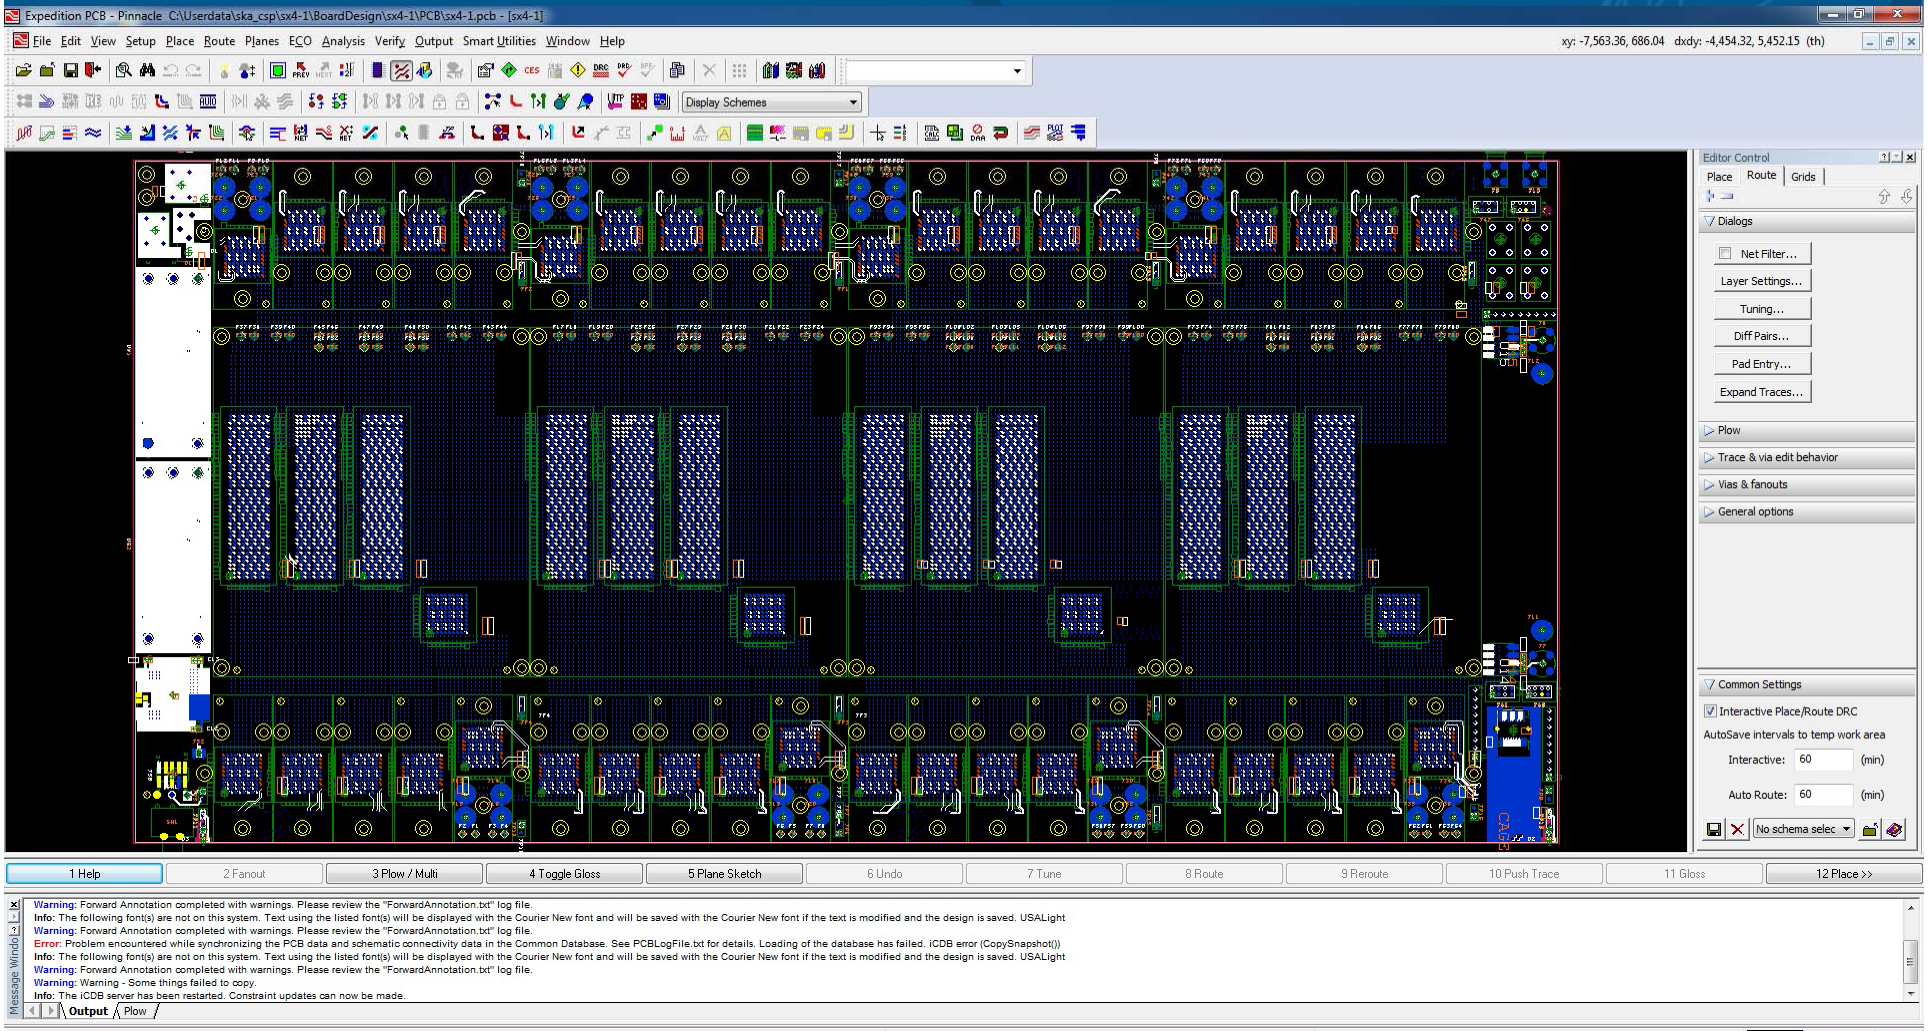

“P32S” PMXM module: Arria/Stratix-10 + 4-link HMC…ready for proto fab (designed by CEI).

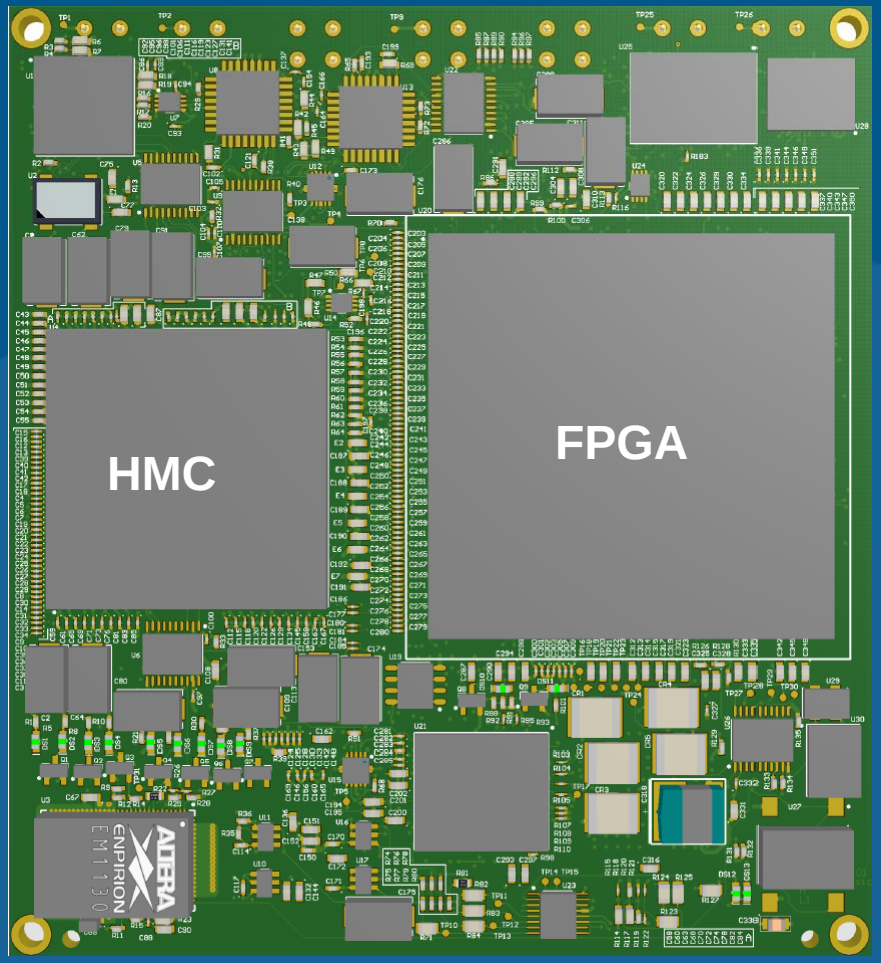

Front Side

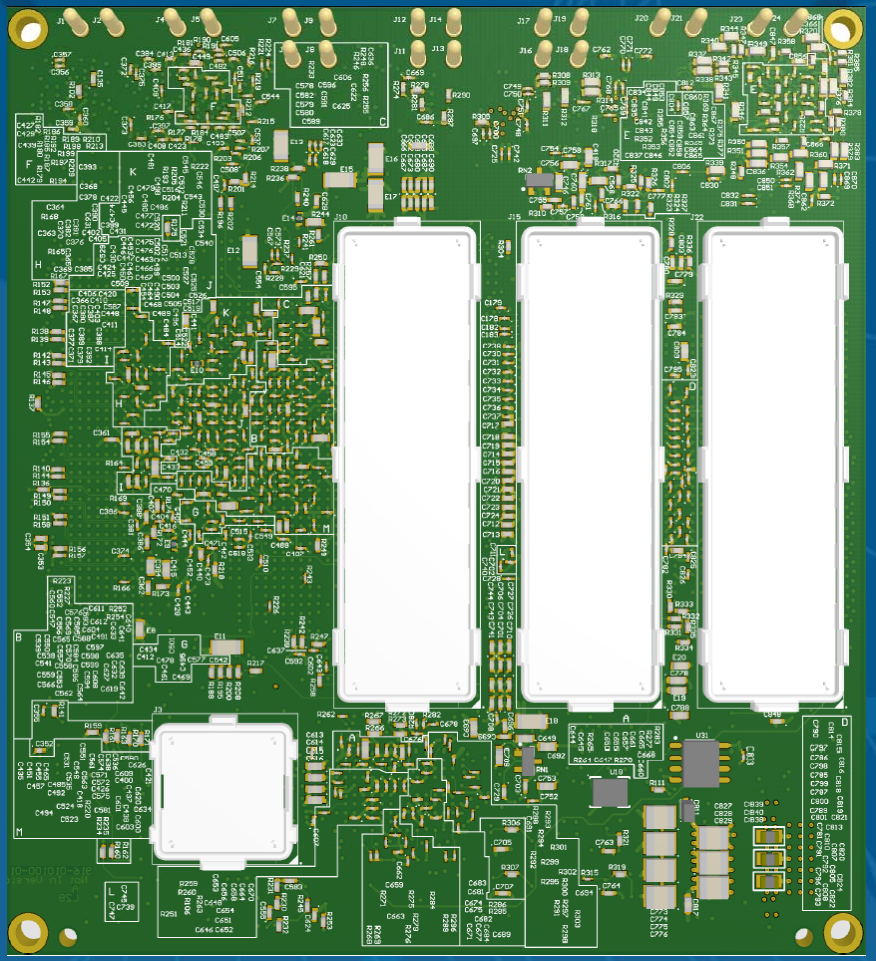

Rear Side

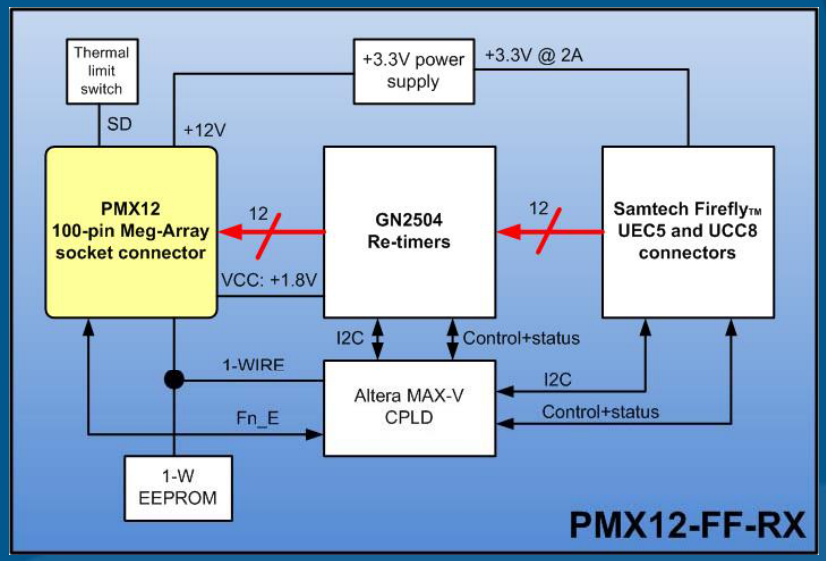

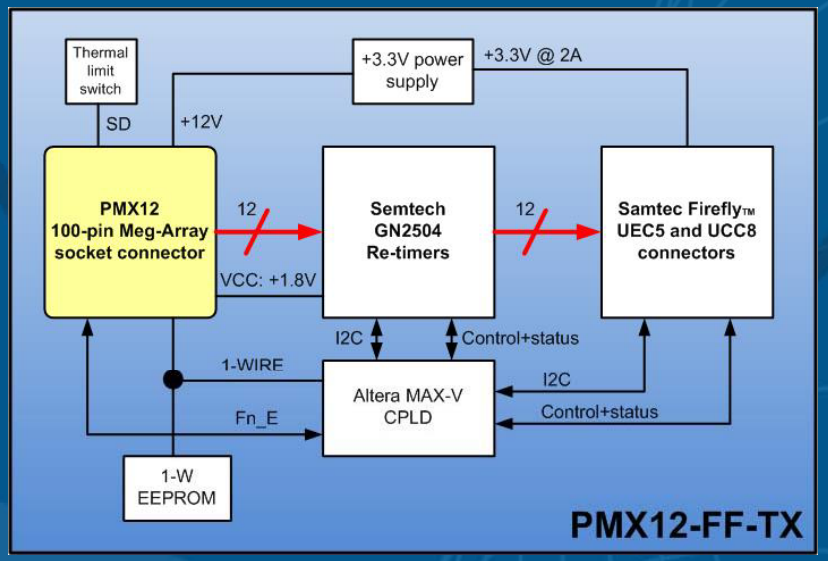

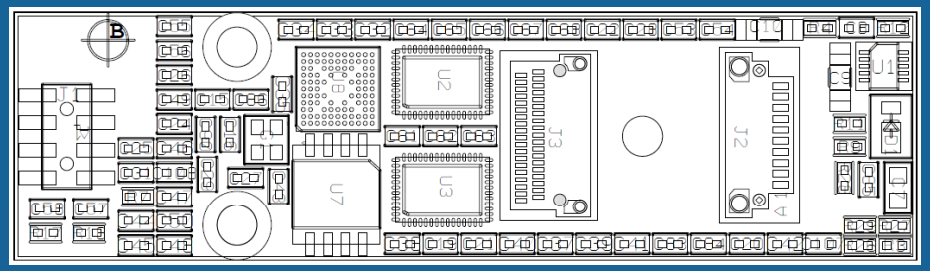

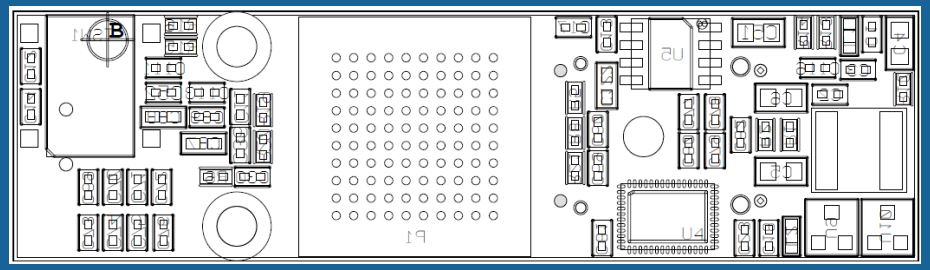

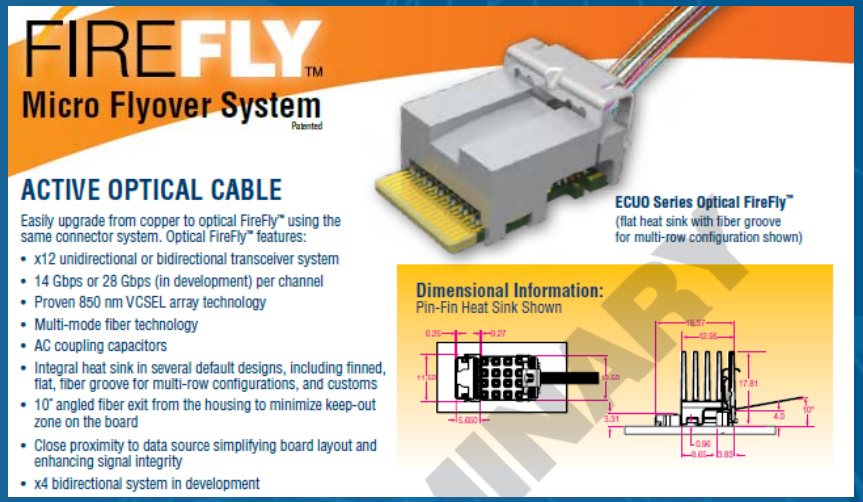

Design is in progress.

On-line order forms/configurators

- The www.powermx.org web site is a goto site for specs and products:
  - PCBAs.
  - Mechanical bits and pieces.
  - Applications(?)

- Currently two such mechanical items are in the works:
  - Parameterized PMXM **active-fan/passive heatsink**—Radian Heatsinks CA, U.S. Done “on spec”. Ready to order parts.
  - Parameterized **pizza box+blade “configurator”** compliant to PMX.1 and PMX.1.2 specifications—Imagination Machine Works, Kelowna Canada. Also done “on spec”.

- No exclusive suppliers.

REL-A-05 March 13, 2015

PowerMX

PRODUCT BRIEF

Web-Based Enclosure Configurator for PowerMX PMX.1-Compliant MB's utilizing PMX.1.2-Compliant I/O Cards

Imagination Machine Works Ltd. will make available in Q2 2015 a web-based configurator for configuration and purchase of various forms of enclosures for the PowerMX PMX.1-compliant Motherboards utilizing PMX.1.2-compliant I/O Cards.

IMW's enclosure designs will range from a simple 1-4 U (in 1/2U increments) 19" rack mount enclosures through to a fully engineered PowerMX blade, backplane, shelf and rack system with integrated cooling solutions.

IMW's PowerMX Enclosure Configurator will offer the user the ability to plug in their required parameters to arrive at the exact enclosure required to perform the needed task. Pricing and 3D models in generic CAD formats will be available directly from the configurator to allow the user to make quick purchase decisions and design changes as required.

Please refer to online standards PMX.1 and PMX.1.2

Standard 19" Rack Mount Enclosure Configurable Features:

- Various heights in U's available in ½ U increments
- Accept motherboards from 1-4 boards, 1-4 PMXM sites each
- 4 sizes and 2 heights available for mezzanine cards (ref. PMX.1.2 specs.)
- 3 bezel types selectable Type-A, Type-B and Type-C (ref. PMX.1.2 specs.)
- Perforated or solid front and/or rear panels
- Liquid cooled configuration available
- Additional front and rear panel holes
- Finish options
- Custom silk-screening

PowerMX Blade, Backplane, Shelf and Rack System Configurable Features:

- Rack system to accept 1-8 shelves. Max. shelves with 4 PMXM site motherboards will be 2 shelves, each with front and rear blades, smaller motherboard configurations will allow for more dense shelf packaging
- Rack to have top mounted hot-swappable fan pack system
- Shelf to accept side installed, backplanes, midplanes or panels
- Shelf to house -48 VDC power distribution bus bars
- Shelves to accept 6-16 blades
- Blade sizes ranging from 1-4 motherboards
- Motherboard PMXM sites from 1-4
- Blades to be hot-swappable
- Custom silk-screening

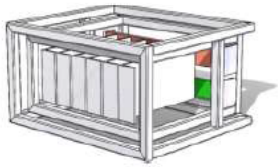

Configurator availability: Q2-2015. For further information email: info@imworks.ca Web: www.imworks.ca

PowerMX is an open-public specification: www.powermx.org.

Imagination Machine Works' pizza-box and blade/shelf/rack configurator.

Range of pizza box sizes, motherboard sites, I/O sizes, heights.

Blade/shelf/rack…single site motherboard, single board to 4 site, 4 board dual blades.

Plug-in parameters, pay $, out come drawings and parts.

All being done "on spec", avail Q2-2015.

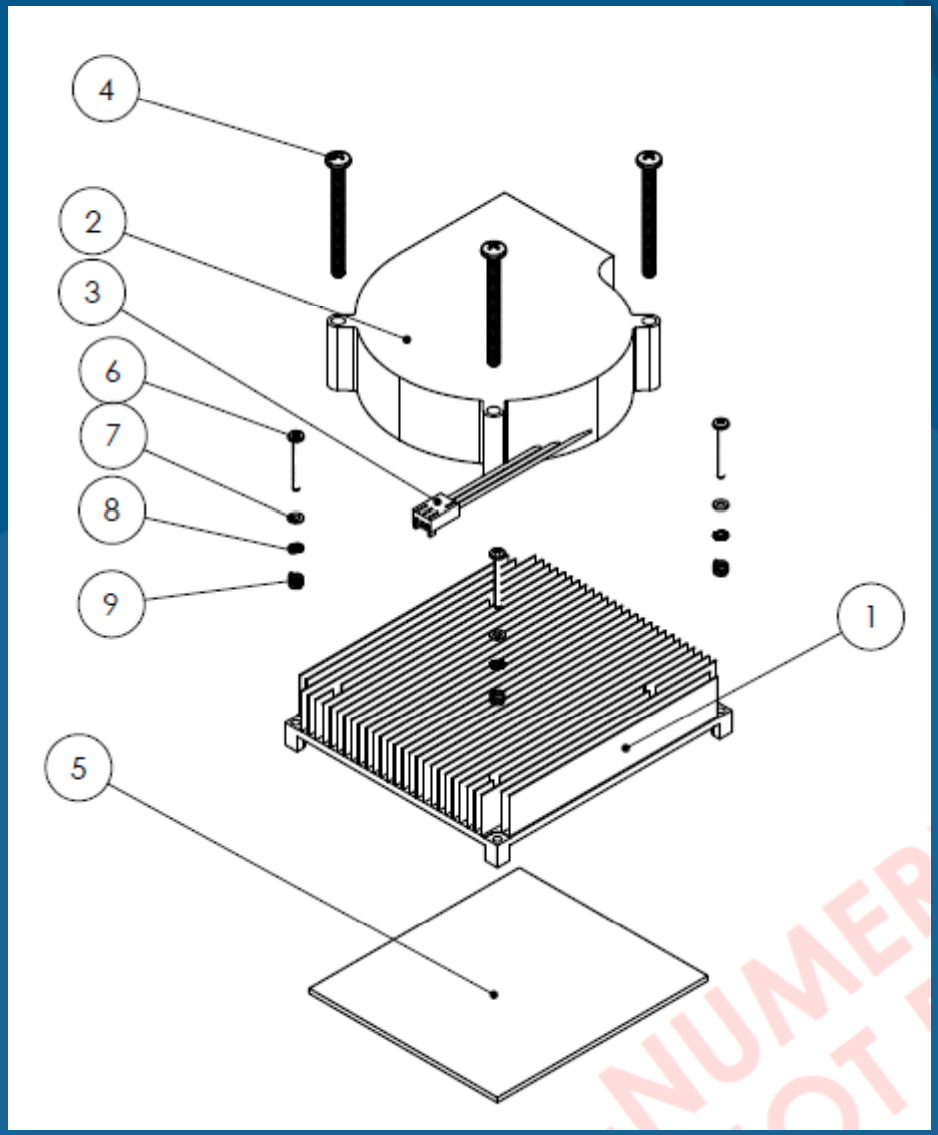

Parameters:

-Fan type (axial, squirrel cage, none).

-Fin height.

-Post height.

-Mounting style (hard or spring).

-Thermal interface material type and thickness.

-Order from Radian Heatsinks…~$125 ea in 100 qty.

ngVLA correlator possibilities

- One possibility: blade/backplane approach based on the NRC MID.CBF design:
  - 256 antennas, 5 GHz/pol'n correlator in a single ~1 m W x ~0.5 m H x ~2 m D shelf.
  - Two shelves in a rack…entire correlator/beamformer in one rack…a substantial system in a ~1 m x ~2 m footprint (~96 kW).
  - 28G SERDES interconnects, 14 nm tri-gate Altera FPGAs.
  - Systolic-array correlation for full band/sub-arraying flexibility…but could be full cross-connect correlator.

Equiv mesh interconnect of uni-dir 512 x 400 GbE

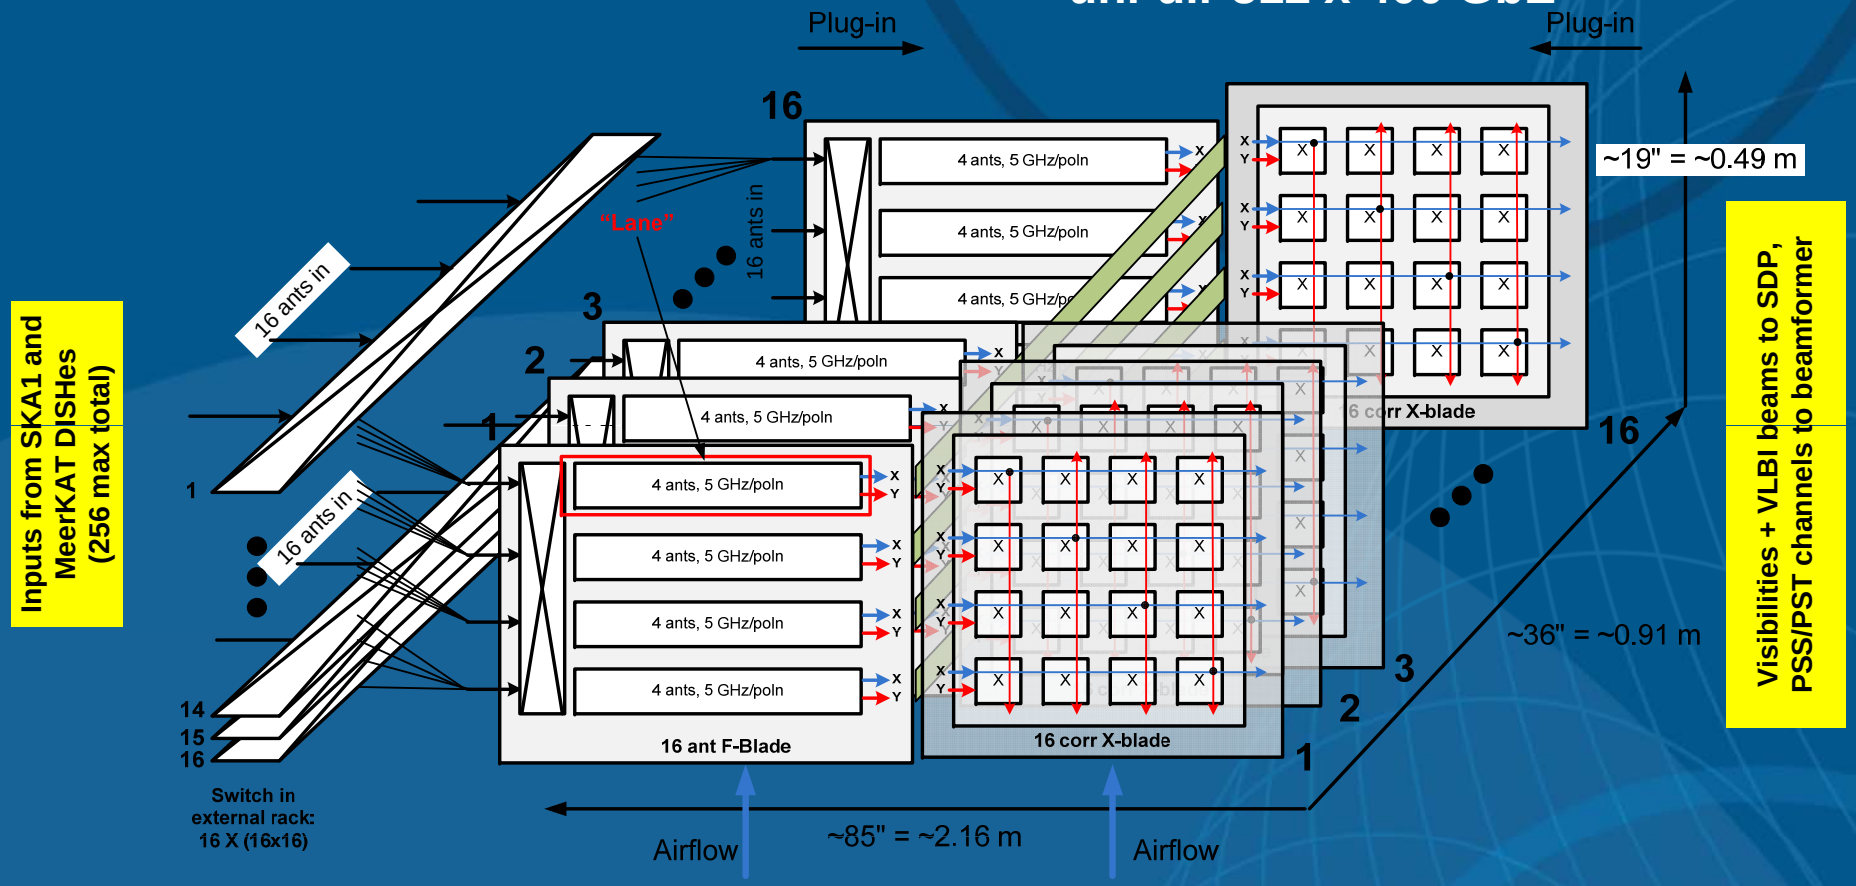

Note: Beamformer in same form factor

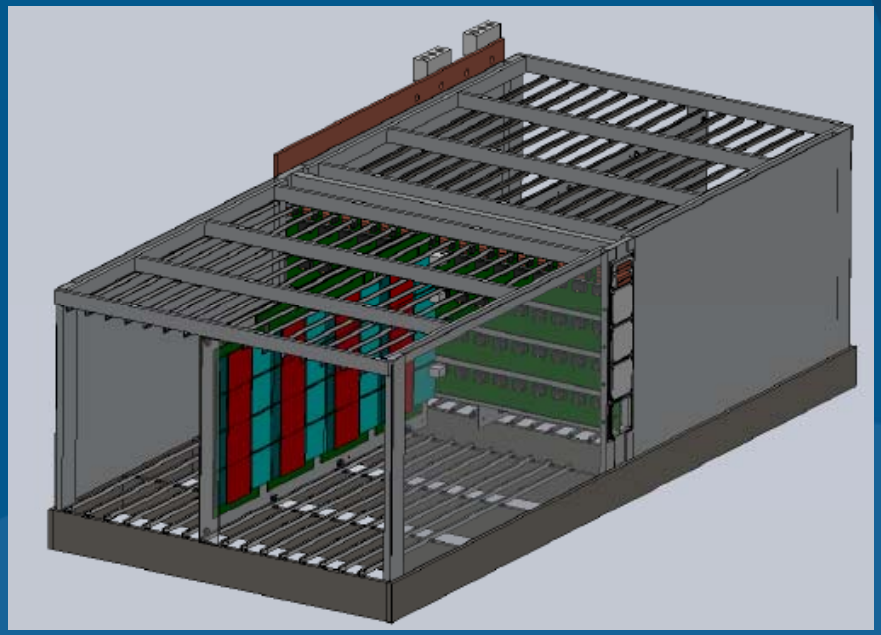

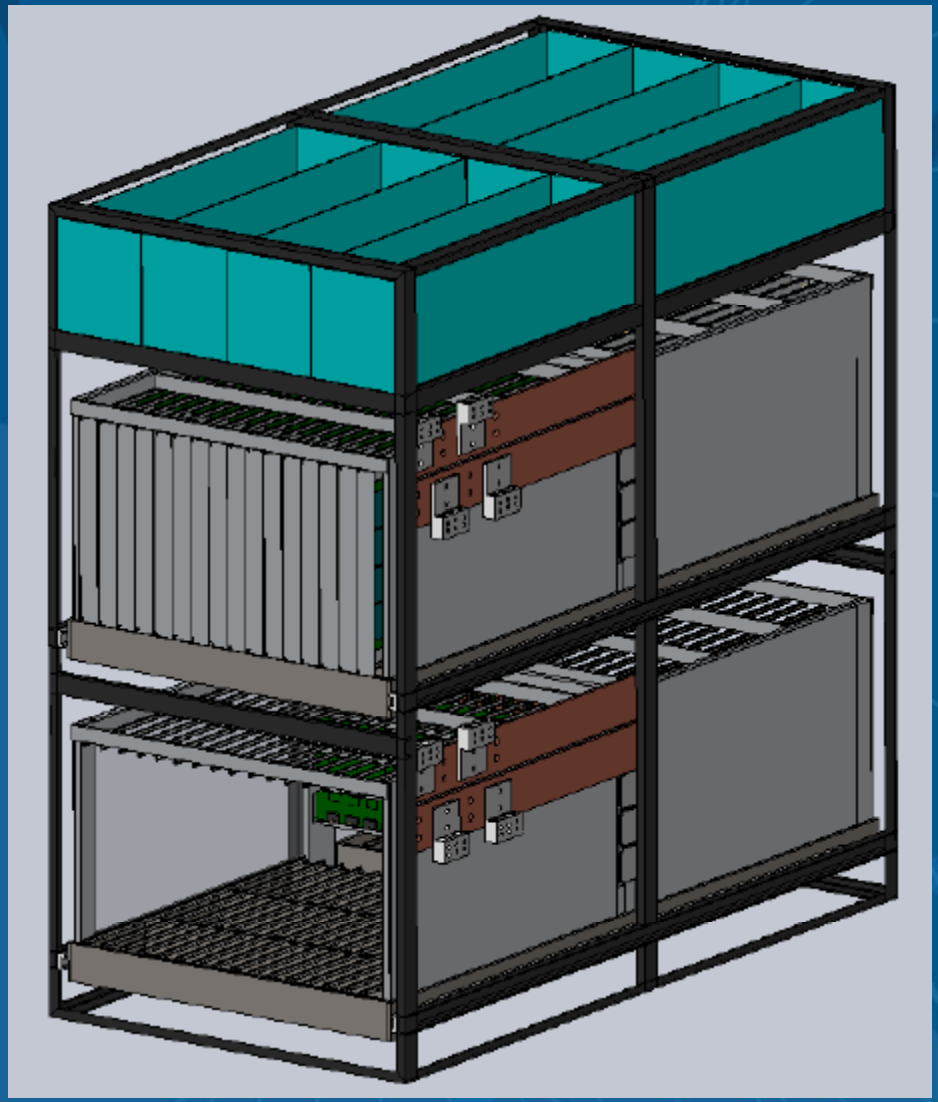

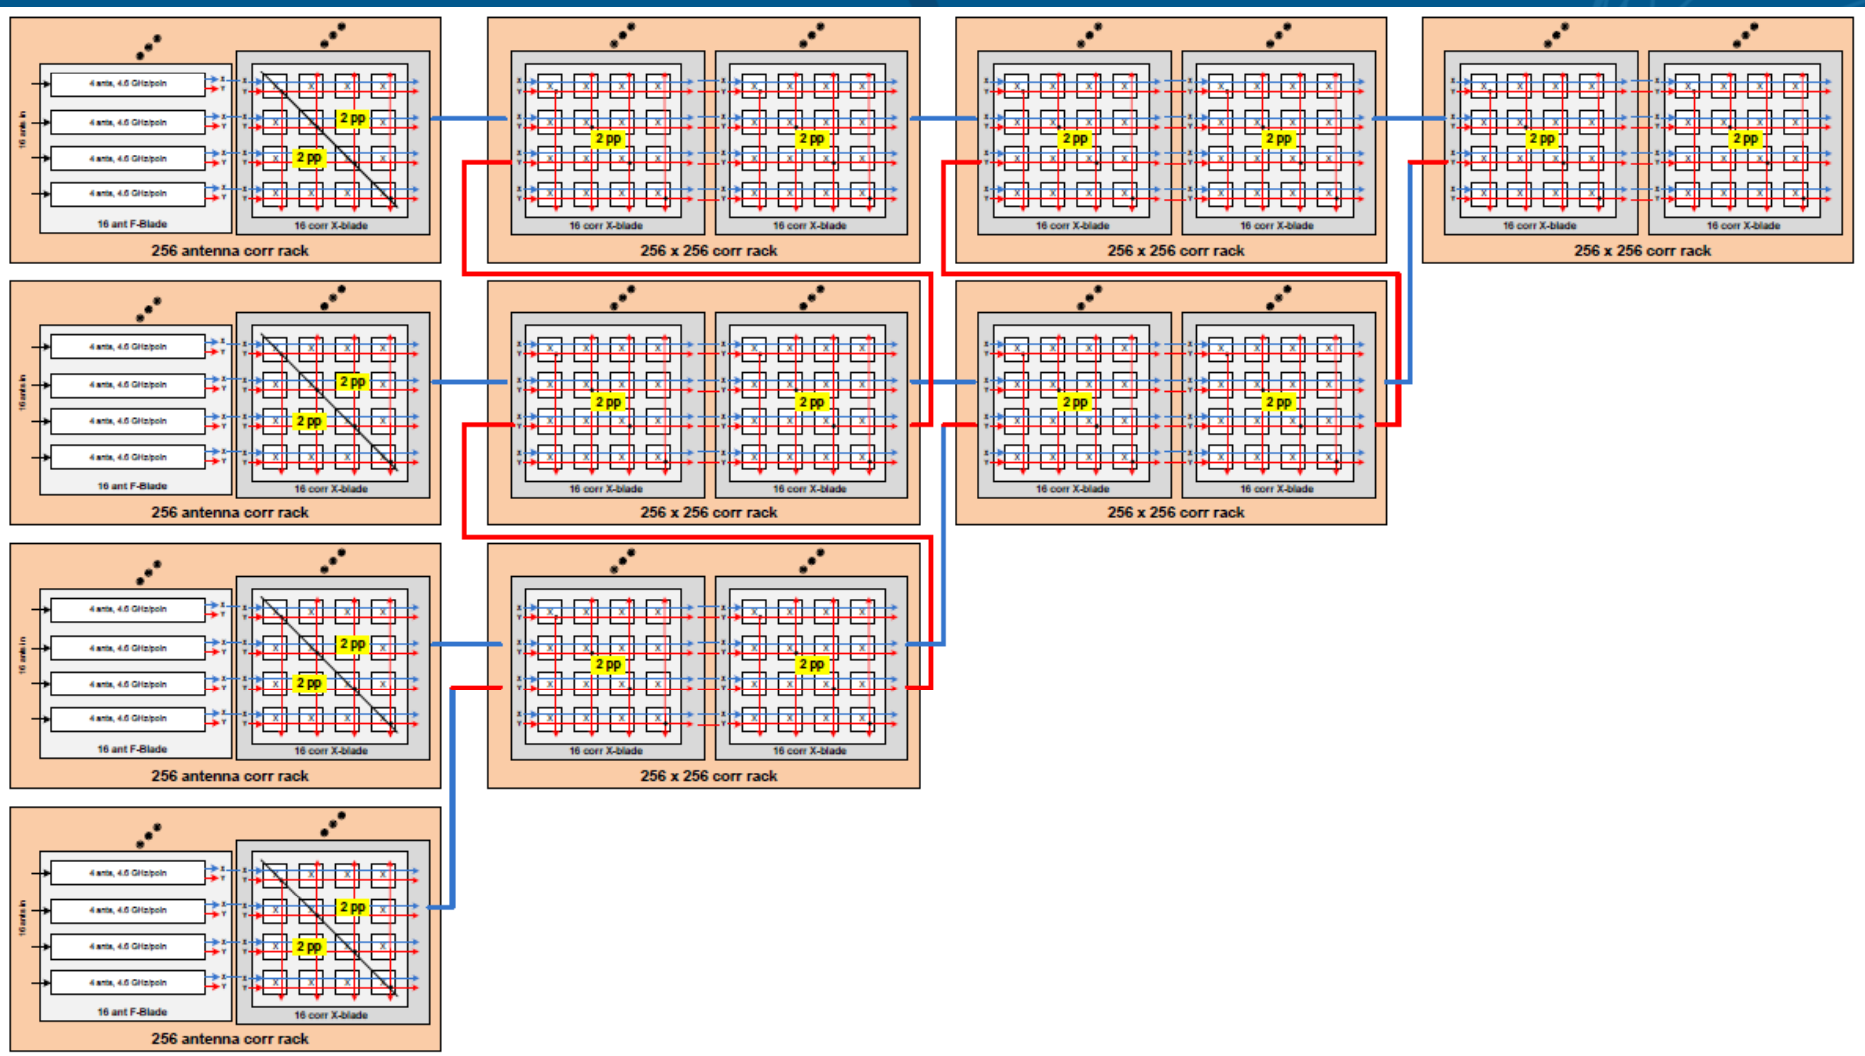

Figure 8-37 Example of scaled blade/backplane correlator to 1024 antennas.

ngVLA correlator possibilities

- Shelf has 10 GHz/pol'n data distribution capability, but likely w/o ASICs, can't process that much BW, even with 10 nm tri-gate (~20-30% performance increase over 14 nm).
- So…
  - 256 antennas, nshelves = ngVLA BW / 5 GHz.
  - 512 antennas, nshelves = 3 x ngVLA BW / 5 GHz.
  - 1024 antennas, nshelves = 10 x ngVLA BW / 5 GHz.

- Or, blade/backplane with fiber mesh interconnect, with liquid cooling for densification:

molex®

Additional Information

VITA 66.1 Ruggedized Optical MT Backplane Interconnect System

Detailed view of installation and maintenance without tools

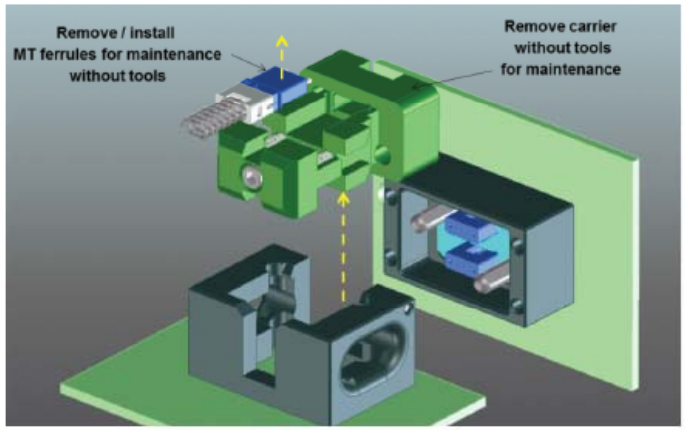

Ease of installation and maintenance

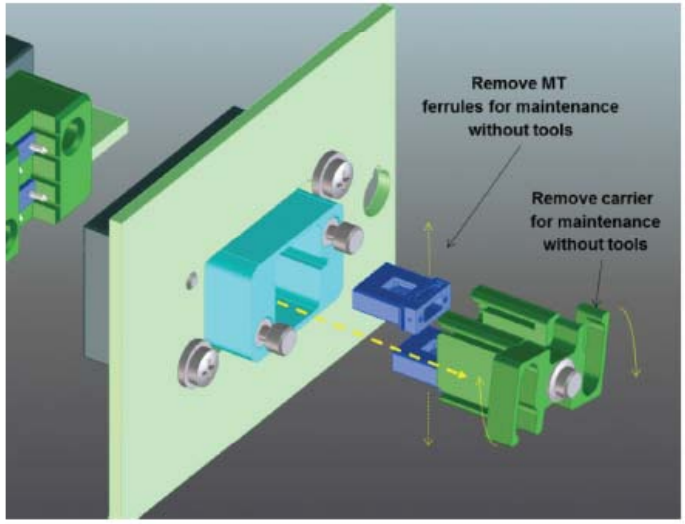

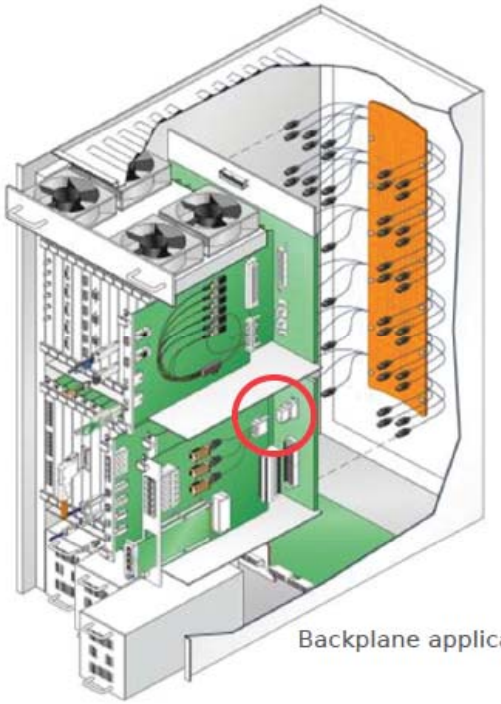

Backplane application

NRC·CNRC

ngVLA correlator possibilities

- Or, pizza boxes with air or liquid cooling, connected by fiber mesh or active switches.

Conclusions

- ngVLA correlator with upcoming SKA-developed technology is feasible.
- PowerMX developments/protos happening now could be used for proto testing and dev (including platform for heterogeneous image processing).
  - Develop PMX-standard but customized cards as required.
- Need to get cost and power down…probably want correlator ASIC.

NRC·CNRC

Thank you

Brent Carlson
Senior Design Engineer
Tel: 250-497-2346
Brent.Carlson@nrc-cnrc.gc.ca
www.nrc-cnrc.gc.ca

Supplementary slides

Overview of re-baselined telescopes

- SKA1-Low:
  - ~128k cross-dipole elements, 300 MHz BW/pol (50-350 MHz), arranged as 512 stations of 256 elements each.
  - Each station requires beamforming.
  - Correlator: 512 stations, 300 MHz, 64k channels/pp.
  - Central Beamforming and Pulsar Search…nbeams TBD, probably <1000.

SKA1-Low: MRO site, Western Australia

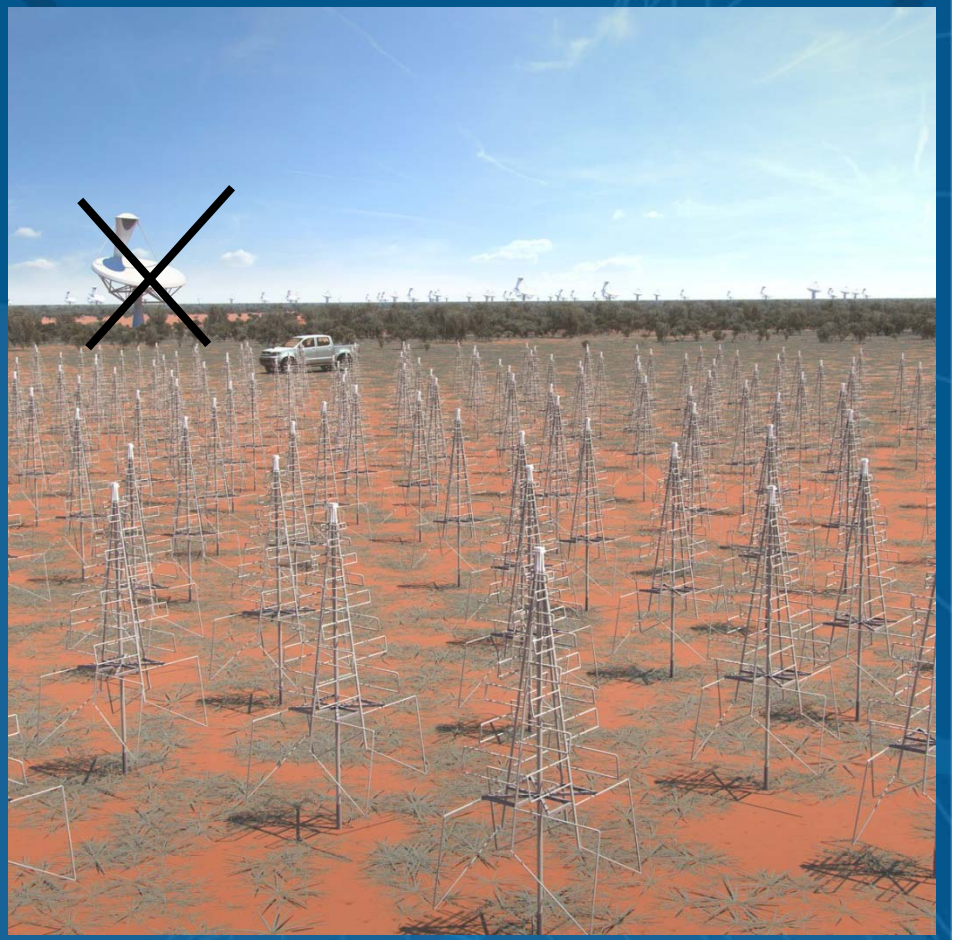

SKA1-Mid: Karoo desert, South Africa

DVA1—full SKA1 Dish prototype, Penticton, BC, Canada

PowerMX

- Motivation:
  - We (engineers in the CSP consortium) are tasked with developing and costing designs now, for production and deployment in the 2018-2022 timeframe.
  - Anything we do now may be obsolete requiring expensive re-design…but limited funds and resources.
  - Is there a way to minimize the cost of the next design cycle?

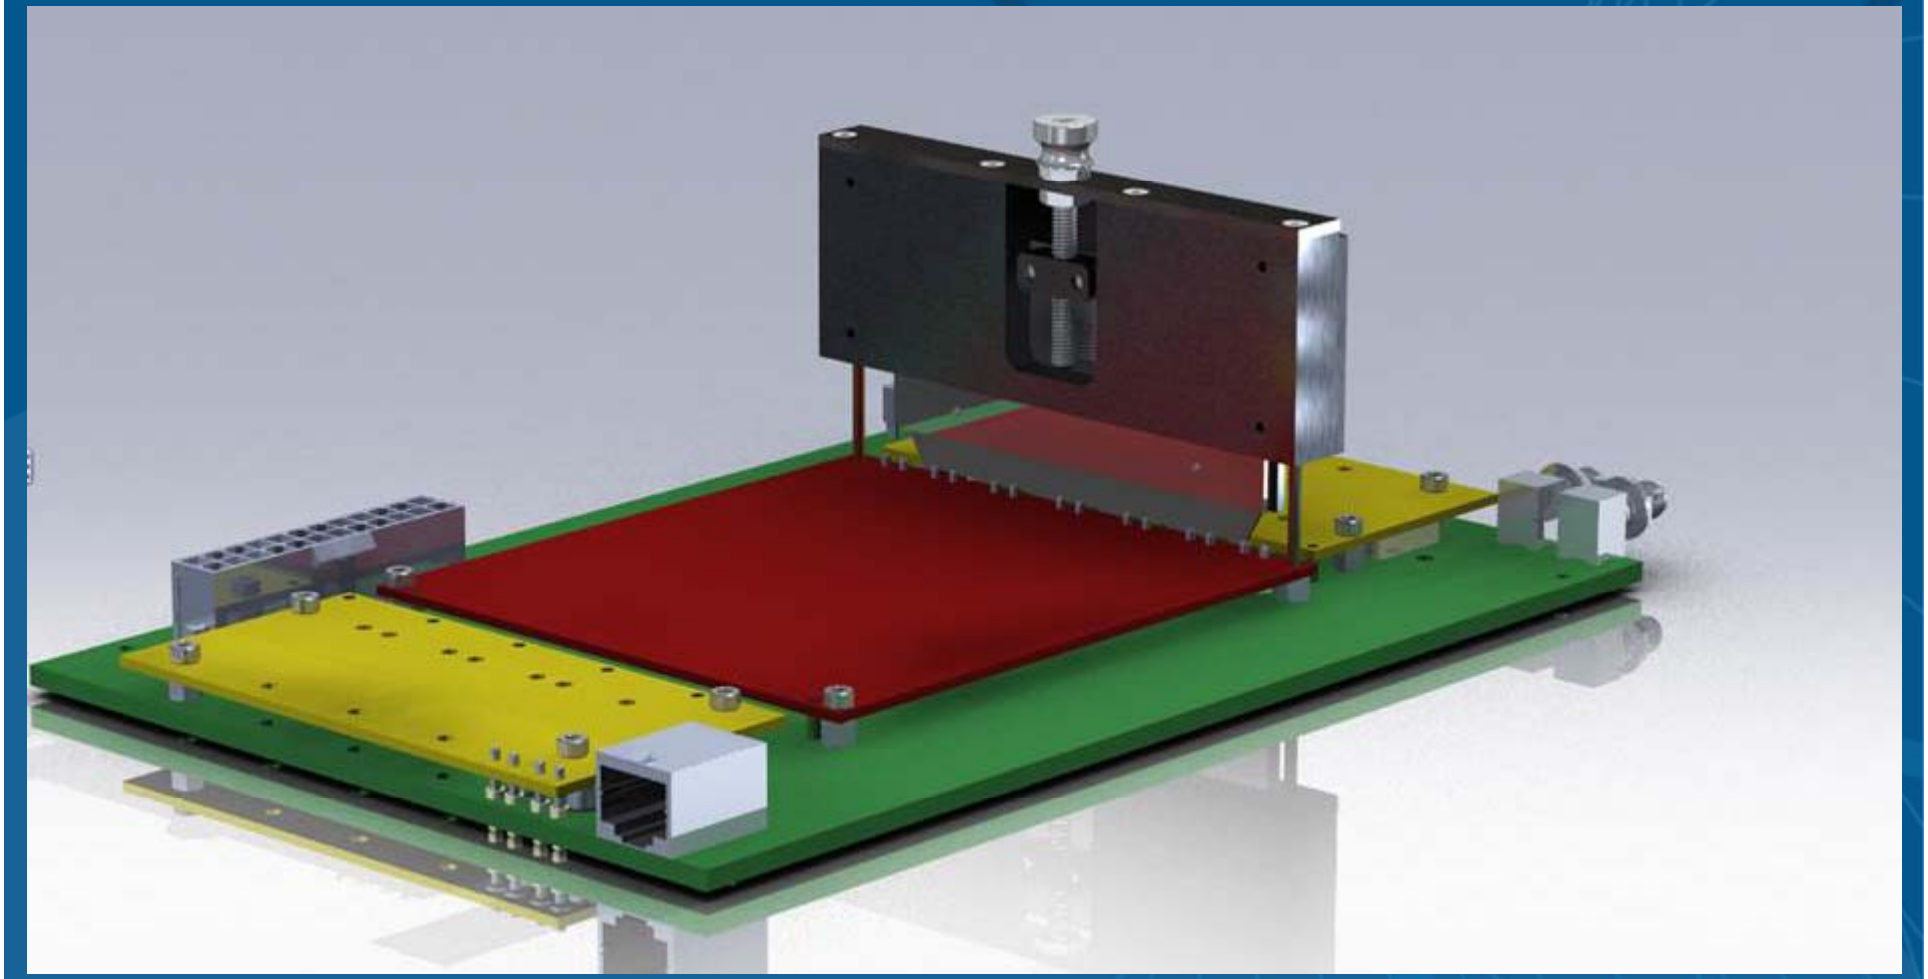

Future work/vision

- Flesh out vision of module plug-n-play compatibility.
  - Further specification layers define module requirements for such compatibility.
  - F/W, host S/W device drivers, “out of the box” GUIs.
  - Definition and set up of repository for applications.

1. Buy modules/motherboards/mechanics for your application.
2. Open box, plug them in.
3. Power up…”self-aware”…download S/W/drivers.
4. Ready for further application development, or ready to start processing.

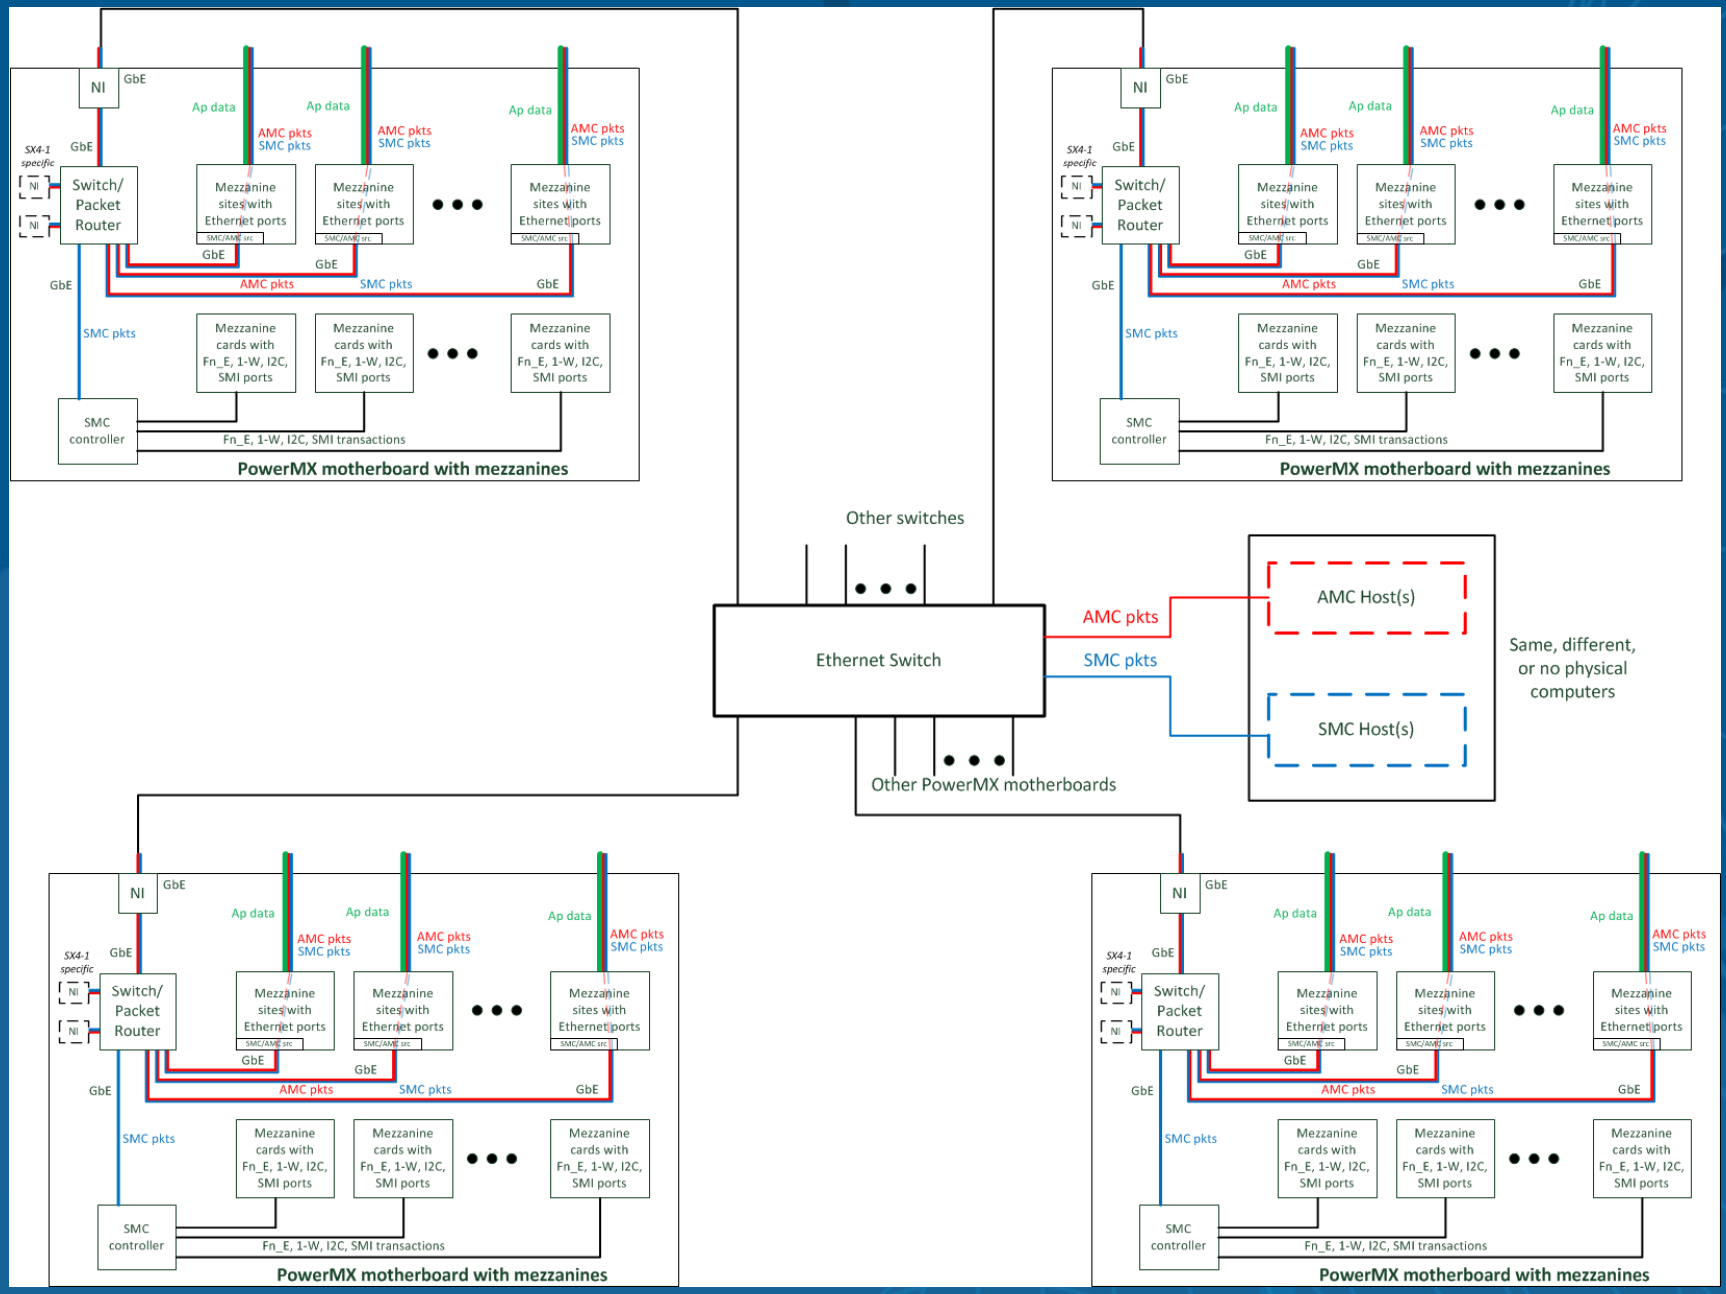
NI
GbE
Ap data
AMC pkts
SMC pkts
SX4-1 specific
Switch/ Packet Router
Mezzanine sites with Ethernet ports
SMC/AMC src
GbE
AMC pkts
SMC pkts
SMC pkts
Mezzanine cards with Fn_E, 1-W, I2C, SMI ports
SMC controller
Fn_E, 1-W, I2C, SMI transactions
PowerMX motherboard with mezzanines
Other switches
Ethernet Switch
AMC pkts
SMC pkts
AMC Host(s)
SMC Host(s)
Same, different, or no physical computers
Other PowerMX motherboards

ngVLA correlator possibilities: next generation blade?

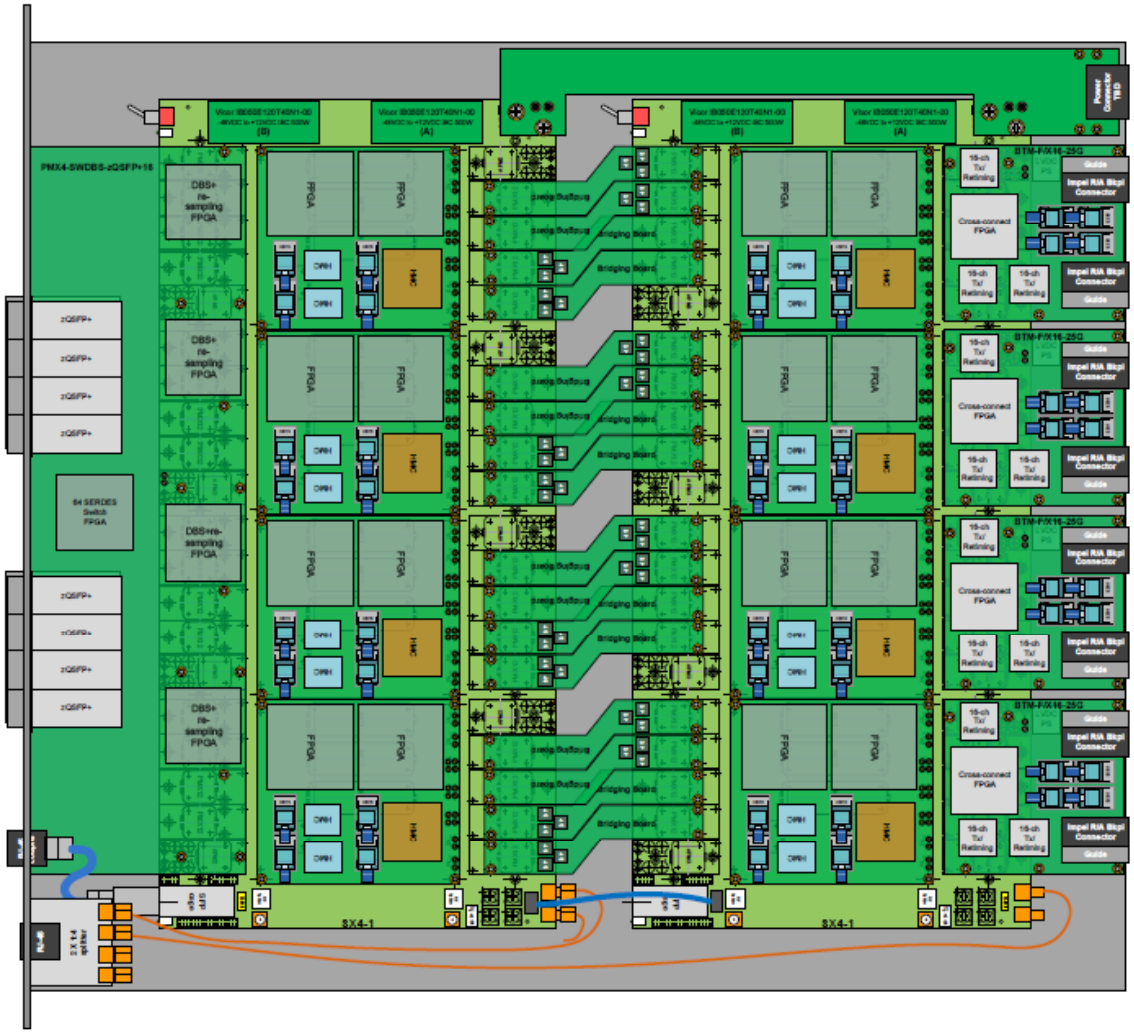

Figure 8-35 2018 technology, notional unified F-X blade. This notion assumes that two 45x45 mm FPGAs and associated HMC and Firefly devices can be integrated onto one PMXM module as in Figure 8-10. Liquid cooling will be required for this blade.

Possibility supported by the IMWorks configurator.