New crate/bus standards: ATCA, uTCA, AMC, VXS and VPX

An overview of the features of the respective standards
and feedback from the IEEE/NSS conference in Dresden

Markus Joos

12 November 2008

Where are we today?

- Accelerator control system uses ~750 VMEbus systems
  - AB/CO looking into uTCA
- DAQ electronics of most LHC experiments based on VMEbus
  - Some xTCA R&D projects started
- DAQ of non-LHC experiments heavily VMEbus based as well

- VMEbus has a number of limitations
  - It is a bus!
  - Transfer bandwidth across backplane (20 – 320? MB/s)
  - Lack of High Availability (HA) features
    - Hot-swap (standardized but rarely used)
    - Redundancy (PSU, cooling, controller)
    - Status monitoring (IPMI standardized but rarely used)
    - Remote control
    - (E)-keying
- But also some strong points
  - Large base of equipment (crates, I/O cards, analyzers, etc.)
  - Lots of experience with H/W and S/W design

What new standards are available?

- PICMG 3.x: ATCA (Advanced Telecommunications Computing Architecture)
- PICMG MTCA.x: MicroTCA/uTCA
- PICMG AMC.x: Advanced Mezzanine Card (for ATCA and uTCA)

- VITA41: VXS
- VITA46: VPX

Not covered in this talk:

- PICMG 1.3: SHB (PCIe based “Industrial PC”)
- PICMG 2.x: Compact PCI (cPCI)
- PICMG EXP.0: PCIe for cPCI
- PCIMG CPLUS.0: A bit like EXP.0 but nothing concrete yet
- PICMG ATCA300.0: ATCA for 300mm deep systems (no rear I/O)

VPX (VITA 46, 150 pages)

Main features:

- Started in ~2003
- Two form factors
  - 3U x 160 mm x 20.32 mm
  - 6U x 160 mm x 20.32 mm
- Supply voltages and power per slot
  - 3U: 3.3V, 5V, 12 V. Max power = 276 W
  - 6U: 5V, 12 or 48 V. Max power = 768 W
  - Cooling (air or conduction) not further defined
- Connector: Tyco 141018(9/7)
- Connectivity
  - 3U: Up to 64 differential pairs @ up to 6.25 GHz (32 if 46.1 implemented)
  - 6U: Up to 192 differential pairs @ up to 6.25 GHz (160 if 46.1 implemented)
  - Switch topology is given by the wiring of the backplane
- Standard only touches HA features
- VITA product directory lists VPX products from 7 companies
- Mainly for (ruggedized) military applications

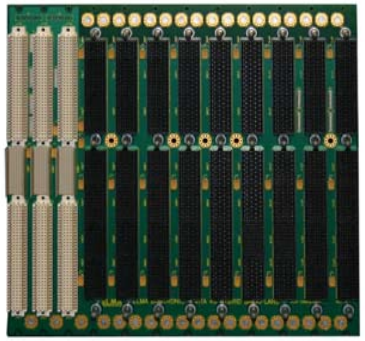

VXS (VITA 41, ~100 pages)

- **Essentially 6U (but 9U not excluded) VMEbus with a new P0 connector**
- **Two types of cards**
  - **Payload**
  - **Switch (one card required, second for redundancy)**
- **Network topology: (dual) star.**
- **Connector: Tyco 1410147-1**
- **Connectivity for payload cards**
  - **16 differential pairs (10 GHz) defined by standard (and routed to switch cards)**
  - **31 reserved pins available on P0**
- **Sub-standards**
  - **41.1: Infiniband**
  - **41.2: Serial RapidIO**
  - **41.3 (draft): IEEE Std 802.3 (1000 Mb/s Ethernet)**
  - **41.4 (draft): PCIe**
- **Hot Swap: According to VITA 1.4**
- **System management based on I2C / IPMI but only formulated as recommendation**

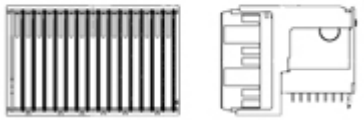

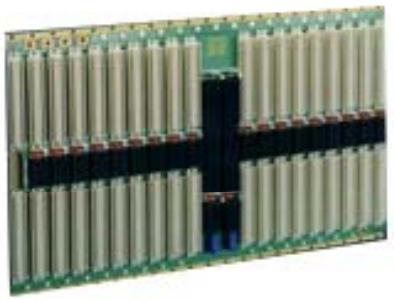

Advanced TCA (650 pages + IPMI)

- More of a system than a board standard
- Started in 2001 by ~100 companies
- One form factor
  - Front: 8U x 280 mm x 30.48 mm (14 slots per 19” crate)
  - Rear: 8U x 60 mm (5W)
- Supply voltage: -48 V (-> DC-DC conversion each on-board)
- Power limit: 200 W (400 W) per card
- Connectors
  - Zone 1: One connector for power & system management
  - Zone 2: One to five ZD connectors for data transfer
  - Zone 3: User defined connector for rear I/O
- Connectivity
  - Up to 200 differential pairs
  - 4 groups
    - 64 pairs for Base Interface (usually Eth., star topology)
    - 120 pairs for Fabric Interface (star or full mesh)
      - Ethernet, PCIs, Infiniband, serial RapidIO, StarFabric
    - 6 pairs for Clock Synchronization
    - 10 pairs for Update Channel
- System management based on IPMI, I2C and FRU data

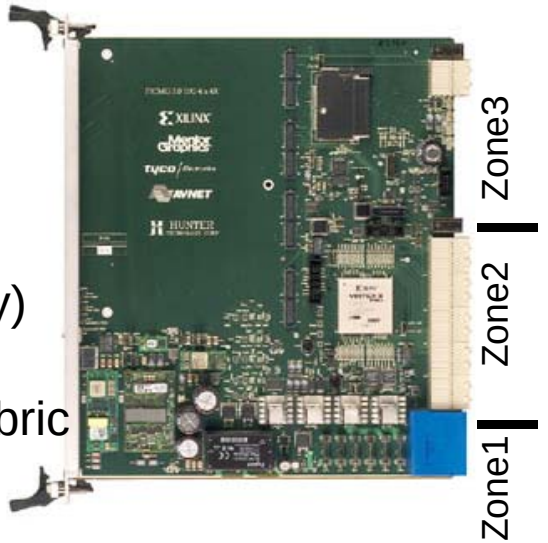

ATCA HA features

(applies also largely to uTCA)

Redundancy

- **Power Supply modules**
- **Ventilators**
- **Shelf managers**
- **Switch blades**

Electronic Keying

- **Based on FRU information payload cards may be accepted / rejected in a given slot**

Hot swap

- **Payload board will only receive (payload) power if the shelf manager can guaranty for the availability of the required resources (power, cooling, signal connections)**

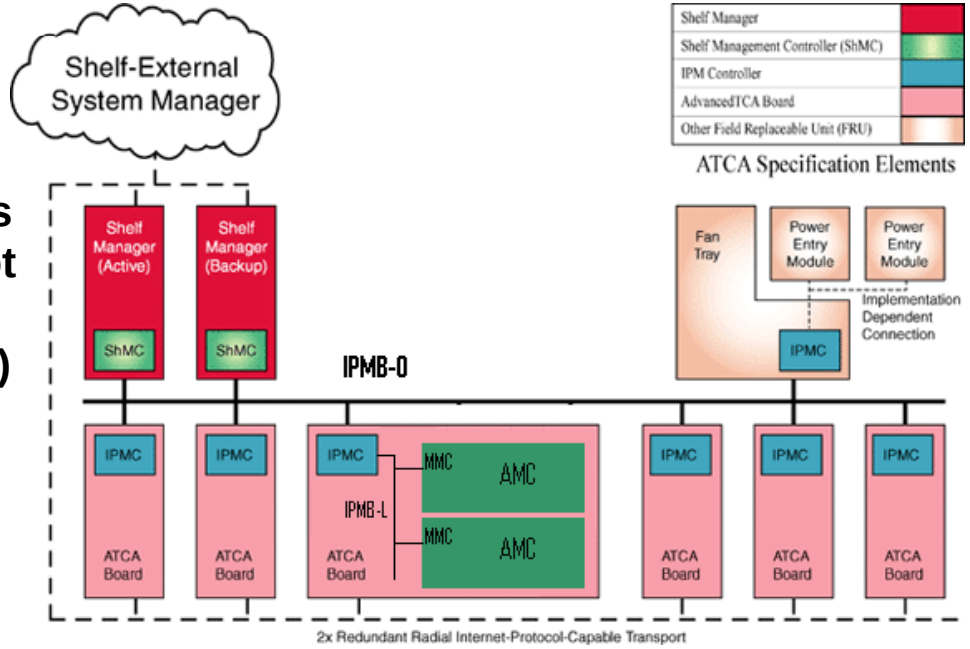

Monitoring

- **Low level: IPMI on I2C**
- **High level: SNMP (Simple Network Management Protocol) and other protocols on TCP/IP**
- **System event logs**

Cooling

- **Dynamically controlled fans and several alarm levels**

Some ATCA equipment

Shelfs

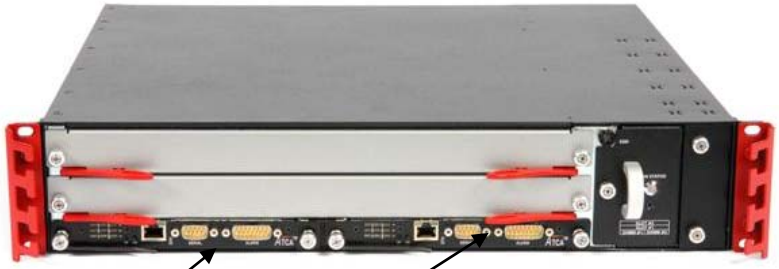

Shelf manager(s)

Hot-swap fans

AMC Carriers

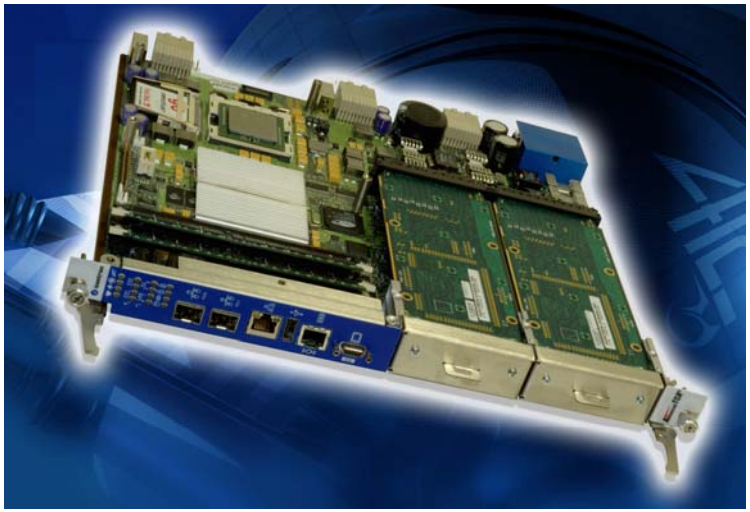

Cutaway carrier (for full-size AMCs)

RTM

Feedback from IEEE/NSS

ATCA trends:

- Backplanes with dual star architecture for the fabric slots more popular than full mesh (Emerson)
- Fabric ports usually based on 1GB or 10 GB Ethernet (Emerson)
- StarFabric dead (Emerson)
- Infiniband, PCIe and RapidIO rarely used (Emerson)
- PCIe not exotic, switch blades exist (powerBridge)
    - As PCIe needs a controller (configuration, interrupt handling) it does not match the ATCA "philosophy"
- AMC carriers are usually of the conventional type. Problem with cut-out carriers is limited PCB space for on-board functions (Emerson)

ATCA observations:

- Companies do not quote end-to-end latencies for e.g. PCIe. Users also don't seem to measure this systematically
    - Eth latency across switch: 250 ns (SLAC)
- RTMs not intended for active electronics (too little power, problems with hot-swap in case of H/W problems)
- A VMEbus adapter for ATCA (company: SAIC) is still under development

AMC

- Originally intended as hot-swappable mezzanine standard for ATCA but soon used as the basis for the uTCA standard
- 6 form factors:
  - 74 or 149 mm wide
  - 13, 18 or 28 mm high
  - 180 mm deep
- Power supply: 80W (max) on +12V (and 0.5W on 3.3V management power)
- Connector: 85 pin (single sided) or 170 pin (double sided) edge connector
- Connectivity
  - Up to 12.5 Gbps
  - 20+20 LVDS signal pairs for data transfer (Eth, PCIe, SAS/SATA, RapidIO)
  - Clock interface, JTAG, I2C (IPMI)

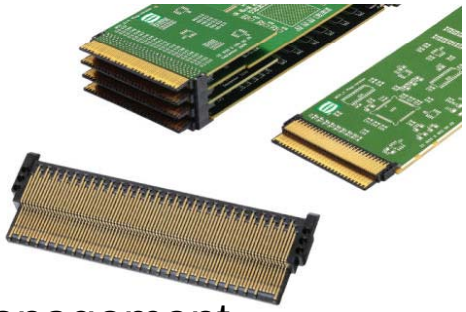

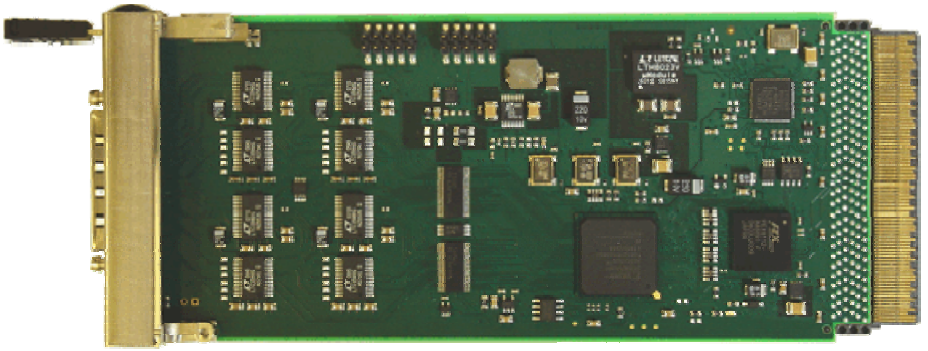

Feedback from IEEE/NSS

AMC trends:

- PCIe, Eth and SATA frequently used (RapidIO exotic)
- No plans for further dot standards
- Compact modules not popular (PowerBridge)

AMC performance:

- As for ATCA little information from the H/W vendors
- 370 MB/s (end-to-end) for x4 PCIe on AMC (DESY)
- PCIe end-to-end latency: 150 ns (DESY)
- PCIe: 200ns write, 2 us read (DESY)

uTCA

- A system standard based on the AMC, standardized in 2006
- Min. signaling speed: 3.125 GHz
- Connectivity:
  - 4 AMC LVDS pairs defined as “Common Options” (2 Eth. & 2 SAS ports) and connect to 1 or 2 MCH boards which provide the switching
  - 8 AMC LVDS pairs defined as (extended) fat pipes (1 or 10 G Eth, PCIe, RapidI/O). Connection to MCH not standardized
  - Remaining 8 LVDS pairs not defined (can be used for rear I/O (but rear I/O not foreseen in uTCA standard))
- System management based on IPMI / I2C
- Hot-swap support for PSU & cooling
- Redundant MCH
- The MCH connector supports up to 84 differential pairs. Therefore only 7 pairs per AMC (based on a 12-slot backplane) can be routed to the switch.

Connector Region	AMC Port #	Signal Conventions				MCH Fabric #
Common Options	0	AMC.2 1000Base-BX				A
	1	AMC.2 1000Base-BX				2/A
	2	AMC.3 SAS				B
	3	AMC.3 SAS				2/B
Fat Pipes	4	AMC.1 x4 PCI-E	AMC.4 x4 SRIO	AMC.2 1000Base-BX	AMC.2 10GBase-BX4	D
	5			AMC.2 1000Base-BX		E
	6			AMC.2 1000Base-BX		F
	7			AMC.2 1000Base-BX		G
Extended Fat Pipes	8		AMC.4 x4 SRIO	AMC.2 1000Base-BX	AMC.2 10GBase-BX4	2/D
	9			AMC.2 1000Base-BX		2/E
	10			AMC.2 1000Base-BX		2/F
	11			AMC.2 1000Base-BX		2/G

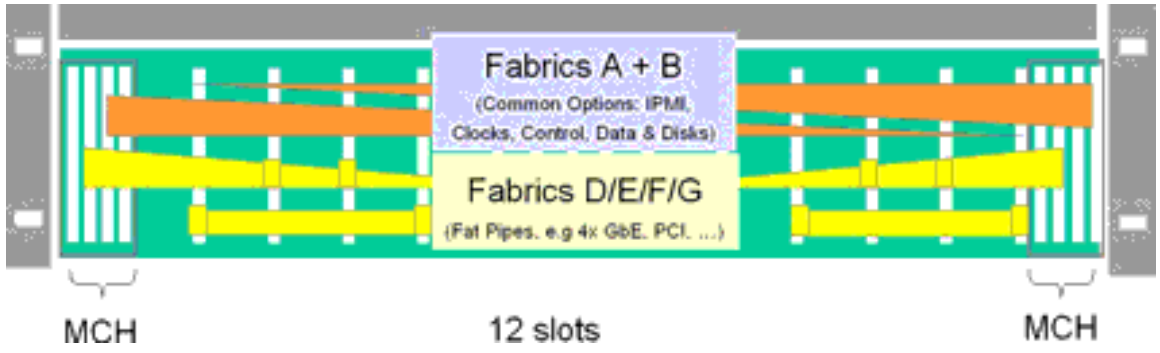

Some uTCA products

CorEdge
- 19” rack mountable
- Dual star backplane
- Up to 10 AMCs
- External AC->DC PSU required

ELMA
- 19” rack mountable
- 8 full size and 4 compact size AMC slots
- For 3rd party power modules

Carlo Gavazzi
- 2U / 19” chassis
- Slots for up to 12 AMCs
- Cooling dimensioned for 40W per slot

Kontron
- 6 mid size (single or double width) AMCs
- AC or DC PSU
- Single star backplane

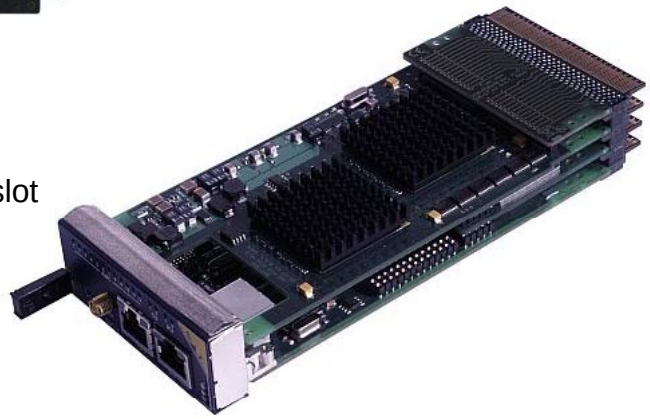

MCH from N.A.T
- Fat-pipe mezzanines for:
  - PCIe
  - 10GB-Eth
  - Serial RapidIO
  - Clocks

Feedback from xTCA workshop @ IEEE/NSS

- Little info on:
  - Switch blades for ATCA
  - Power supply solutions (AC, DC) for ATCA
  - Use of COTS CPUs
  - Backplane topologies
- Manufacturer presentations focussed on:
  - System management (hot-swap protocol, thermal management)
  - Standard (open source) S/W for system management
- (Some) current xTCA projects
  - DESY: XFEL -> uTCA and ATCA for machine (LLRF) control
  - GSI: PANDA -> general purpose ATCA processor with 4+1 FPGAs and fiber I/O
  - Several current and future projects in France (LeDu)
  - CMS GCT
- Efforts to provide open source management S/W for xTCA: OpenSAF, SAForum
- Efforts to insure interoperability of xTCA products: CP-TA, SCOPE alliance
- Many vendors seem to be in favour of “profiles” that limit the number of options given by the standards
  - Proposal for a “physics profile” under the umbrella of PICMG-US

xTCA degrees of freedom (not necessarily a complete list)

- ATCA
  - Communication protocol(s) on the fabric channels
  - Routing of the fabric channels on the backplane (network topology)
  - Connection between front board and RTM
  - PCB thickness (1.6 – 2.4 mm ?)
  - Degree of redundancy
  - Power supply at shelf level (230 VAC or -48 VDC)
- AMC
  - Card height (13, 18 & 28 mm)
  - Card width (74 & 149 mm)
  - Communication protocols (currently 4 options)
  - Number of pins on the connector (85 or 170)
  - JTAG support
- uTCA
  - AMC height & width
  - Degree of redundancy (MCH, PSU, cooling)
  - Routing of the fabric channels on the backplane
  - JTAG support
  - Connectivity of MCH to backplane (1 to 4 tongues) and type of communication protocol on the fat pipes

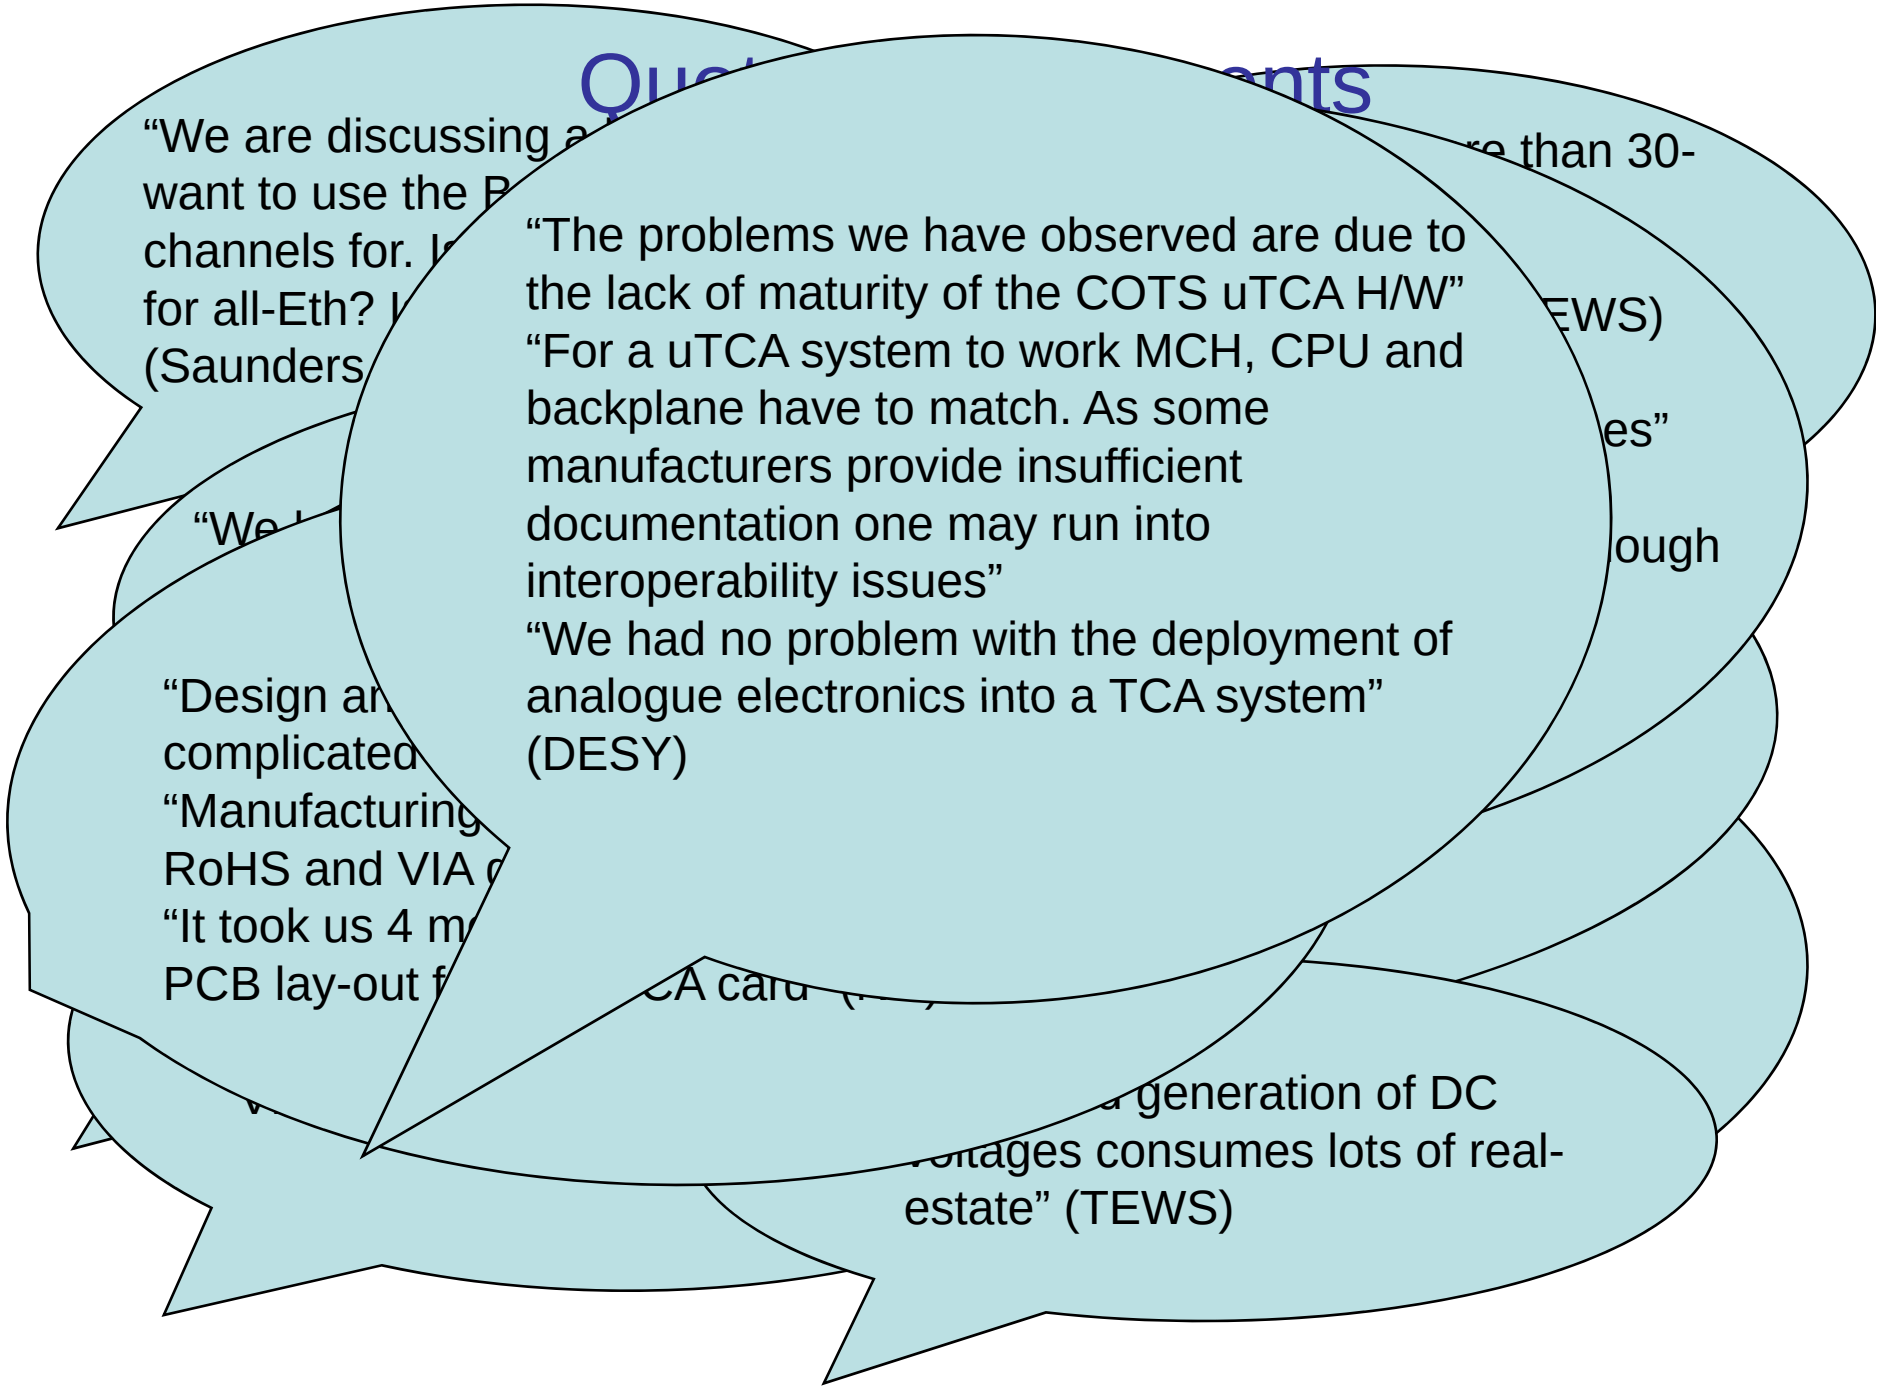
Quest ents
"We are discussing a
want to use the P
channels for. I
for all-Eth? I
(Saunders
re than 30-
EWS)
es"
ough
"We
"Design ar
complicated
"Manufacturing
RoHS and VIA c
"It took us 4 m
PCB lay-out f
CA card"
"The problems we have observed are due to the lack of maturity of the COTS uTCA H/W"
"For a uTCA system to work MCH, CPU and backplane have to match. As some manufacturers provide insufficient documentation one may run into interoperability issues"
"We had no problem with the deployment of analogue electronics into a TCA system" (DESY)
generation of DC
oltages consumes lots of real-estate" (TEWS)

So, what is the best strategy for us?

Predictions are notoriously difficult to make — especially when they concern the future!
(Mark Twain)