

IEEE 802.3 10Mb/s Backplane Ethernet Call For Interest Consensus Presentation

Amrik Bains, Cisco

David Tremblay, HPE

George Zimmerman, CME Consulting

Jon Lewis, Dell EMC

Mandeep Chadha, Microsemi

Orlando, FL

Nov 9th-12th, 2017

Introductions for today's presentation

Presenters and Expert Panel:

Amrik Bains

Cisco

David Tremblay

HPE

George Zimmerman

CME Consulting

Jon Lewis

Dell EMC

Mandeep Chadha

Microsemi

Agenda

- Overview Discussion Jon Lewis, Dell EMC
- Presentations
 - 10BP in Servers Jon Lewis, Dell EMC
 - 10BP in Switches Amrik Bains, Cisco
 - 10BP Technical Feasibility Mandeep Chadha, Microsemi
 - 10BP Why Now? George Zimmerman, CME
- Q&A
- Straw Polls

CFI Objectives

- To gauge the interest in studying 10Mb/s Backplane Ethernet
- We do not need to:
 - Fully explore the problem
 - Debate strengths and weaknesses of solutions
 - Choose a solution
 - Create a PAR or 5 Criteria
 - Create a standard
- Anyone in the room may vote or speak

Overview: About 802.3cg/10SPE

- Single pair connection
- Low cost/data rate suitable for sensors
- Supports power and data
- Supports multiple reaches (e.g. 15m, 1000m, etc.)
- Targets Industrial and Building Automation installed cable

Overview: Intra-system Management Interface

- Internal serial control interfaces are common in server and switches.
- Existing serial interfaces form a high percentage of design/debug/support issues.
- 10SPE could replace a number of these (e.g., I2C/SMBus, MDIO).
- Systems shipping in 10s of millions annually contain high 10s to low 100s of these links.
- Implementations using FPGAs or micro controllers could support a “faster & richer” interface than existing interfaces and be easier to debug using standard Ethernet tools.


10BP in Servers

Jon Lewis, Dell EMC

Server market size info

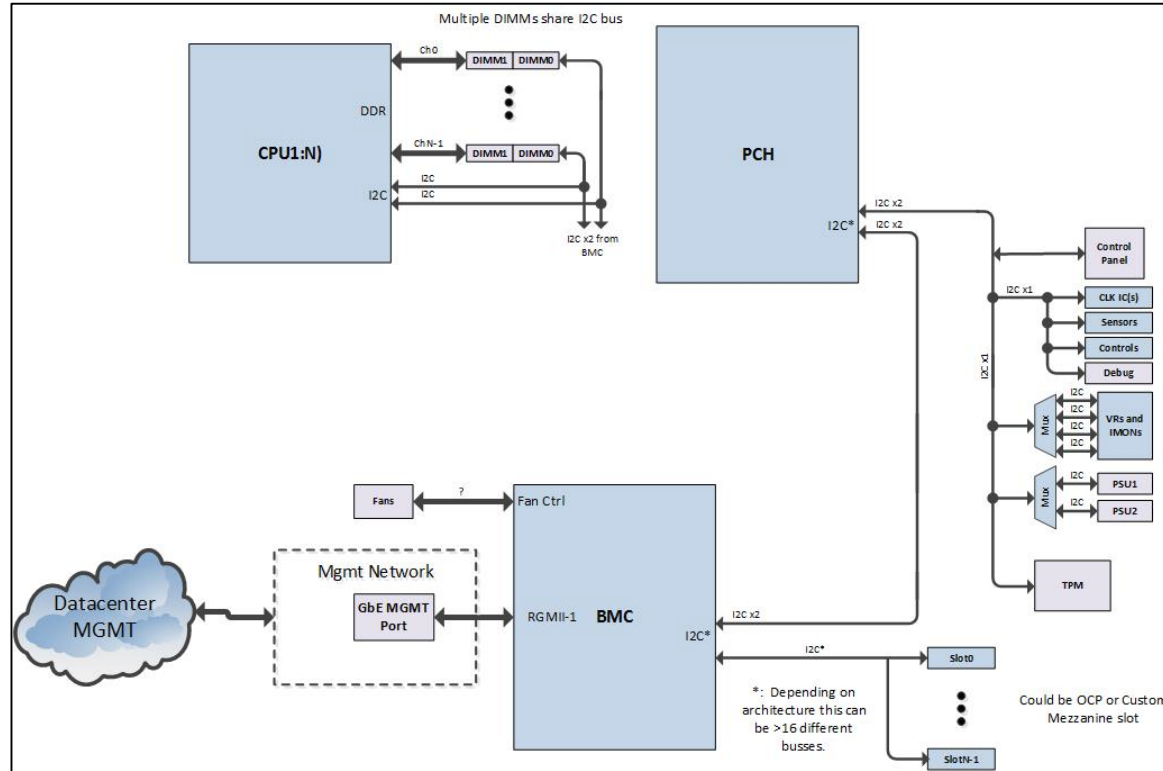
WW 2017Q2 x86 Server Annual Unit Forecast
(units in thousands)

- *Total endpoint device >200 million parts per year*
 - *>20 endpoints per Server*
- *Total switch device >10 million per year*
 - *~1 switch per Server*
- *200 million endpoints*
- *10 million switches*



IDC WW Server Forecast Tracker - September 20, 2017


Current Architecture



Why 10BP?

- Ethernet provides a standard ubiquitous management communication path.
- Same number of pins as existing interfaces maintaining current pin count while increasing functionality.
- Ethernet allows for discovery of devices using a common driver.
 - Going from many custom drivers to a standard Ethernet driver greatly reduces coding complexity and validation time.
- As compute and network nodes “converge” there is a fine line between the server and the network
 - Bus Expanders
 - Chassis servers
 - Modular server implementations.
- Automated alert support is more robust than multi-master SMBus.
 - Using SMBus multiplexers makes multi-master difficult/impossible, requiring continuous bus scanning.

Desired Architecture



10BP in Switches

Amrik Bains, Cisco

Switch market size info

- *Switch Market units per year*
 - *fixed switches* ~x1 million units per year
 - *modular switches* ~y1 million units per year
- *Average number of devices in*
 - *fixed switches* ~x2 endpoints, ~x3 switches
 - *modular switches* ~y2 endpoints, ~y3 switches
- *Total devices per year*
 - *endpoints* >z1 million
 - *switches* >z2 million

Current Switch Architecture

- Internal control plane is used to perform configuration/monitoring of components in the switch
- Many different components with various control plane interfaces, e.g.

➤ I2C/SMB Bus

- Optical Module, AC/DC Power Supplies, FAN Control, DC-DC Converters, Temp monitors, EPROM etc....
- 2 wire – Clock + Shared Data (Tx/Rx)
- Clock speed 100KHz, data BW (25 to 30Kb/s) limited due half-duplex and protocol overhead

➤ UART: Universal Asynchronous Rx/Tx

- Micro-controllers/CPU: Console port
- 2 wire – Rx/Tx
- 9.6Kb/s

➤ MDIO (IEEE 802.3 Clause 22/45): Copper PHYs and Fiber PHYs

- 2 wire
- Tx/Rx shared (half-duplex)
- Max specified MDC clock of 2.5MHz (avg. BW 1Mb/s)

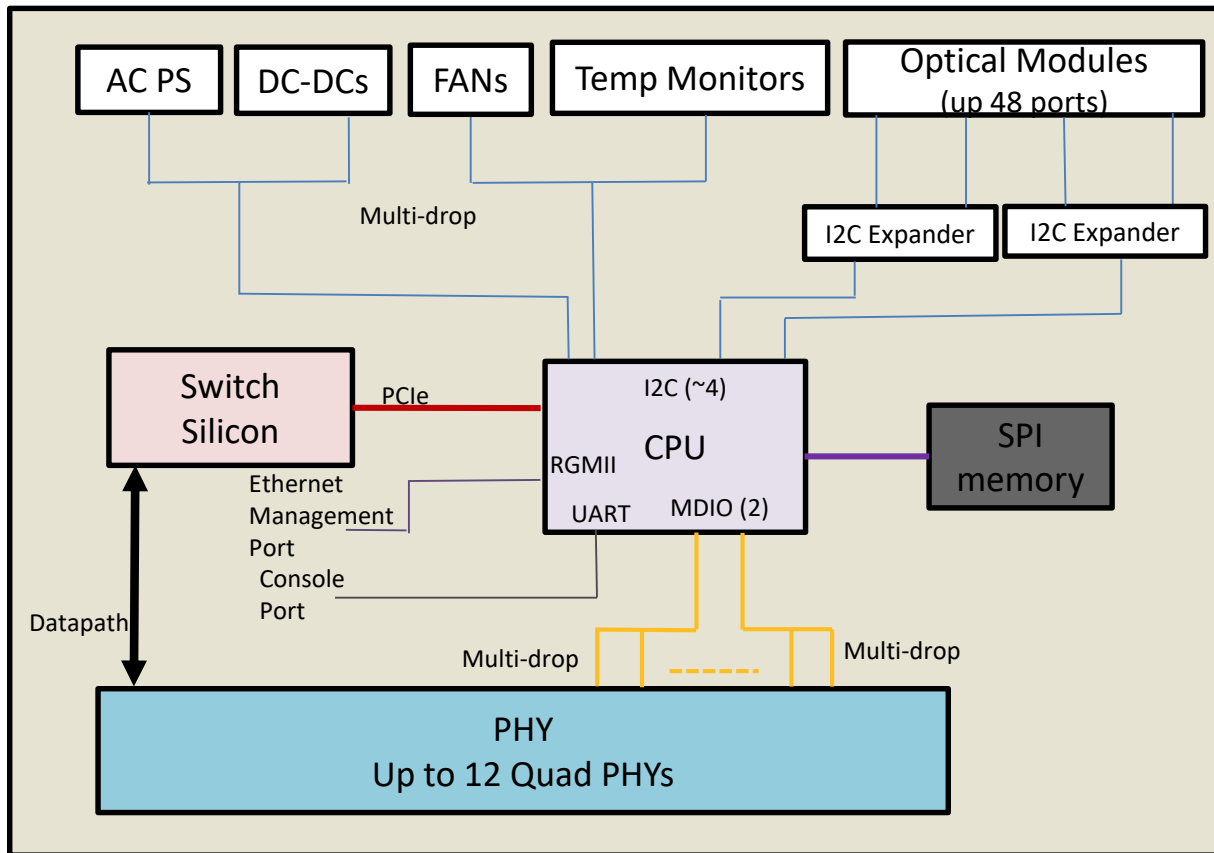
➤ SPI :

- SD card, Sensors, eMMC
- Minimum of 4 wire – Clock, Rx, Tx and Slave Select (incase of multiple devices connected to same data pins)
- Typically 12 to 25Mb/s
- Allows more data wires to be used for higher bandwidth

- Management Interfaces have not kept up with BW requirements
- Different Software Drivers for each interface type

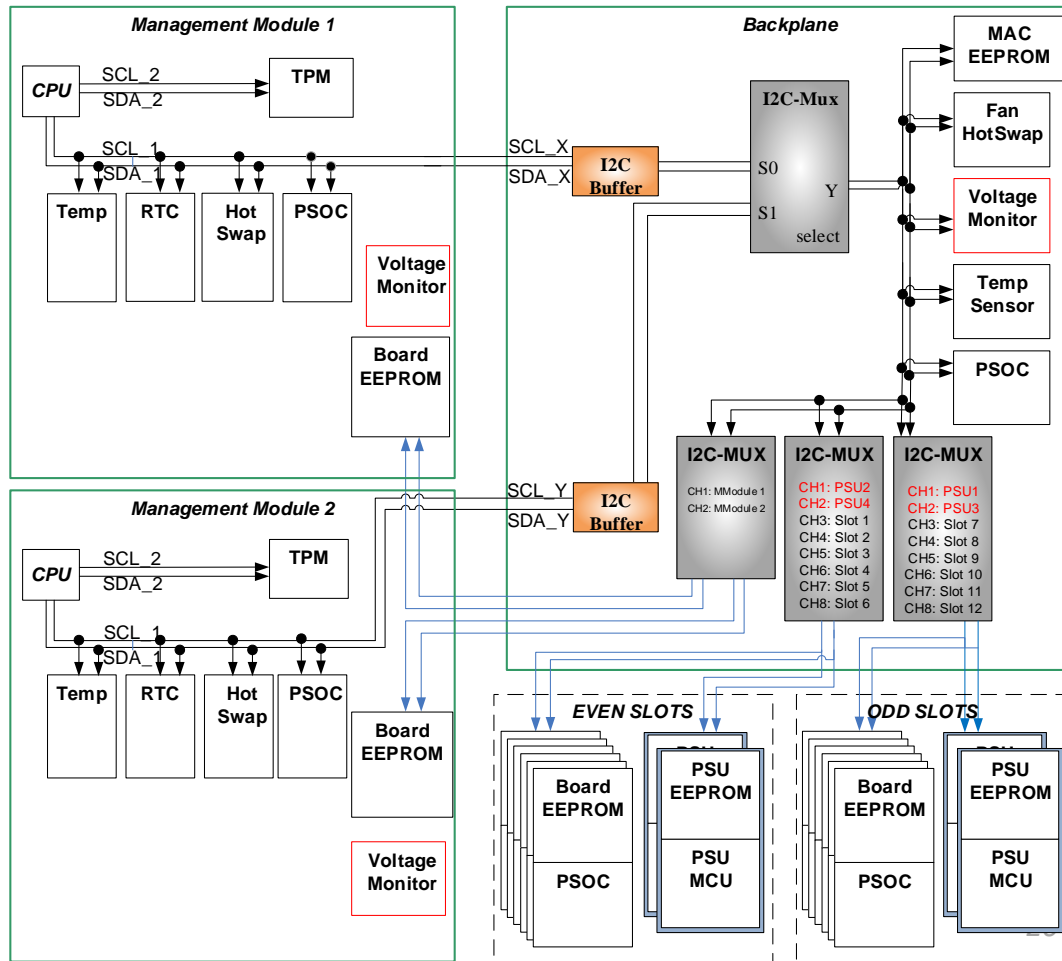
Current Fixed Switch Architecture

- Too many interface types
- I/O types limited on CPU – require external devices to expand I/O
- Multi-drop – subject multiple failure conditions
- Limited BW
- Multiple Software Driver requirements



Current Modular Architecture

- CPU has limited I/O – requires external devices to expand I/O
- Multi-drop – subject to failure
- SW Driver Variation – requires excessive tuning and validation
- Limited Bandwidth



Slide: David Tremblay, HPE

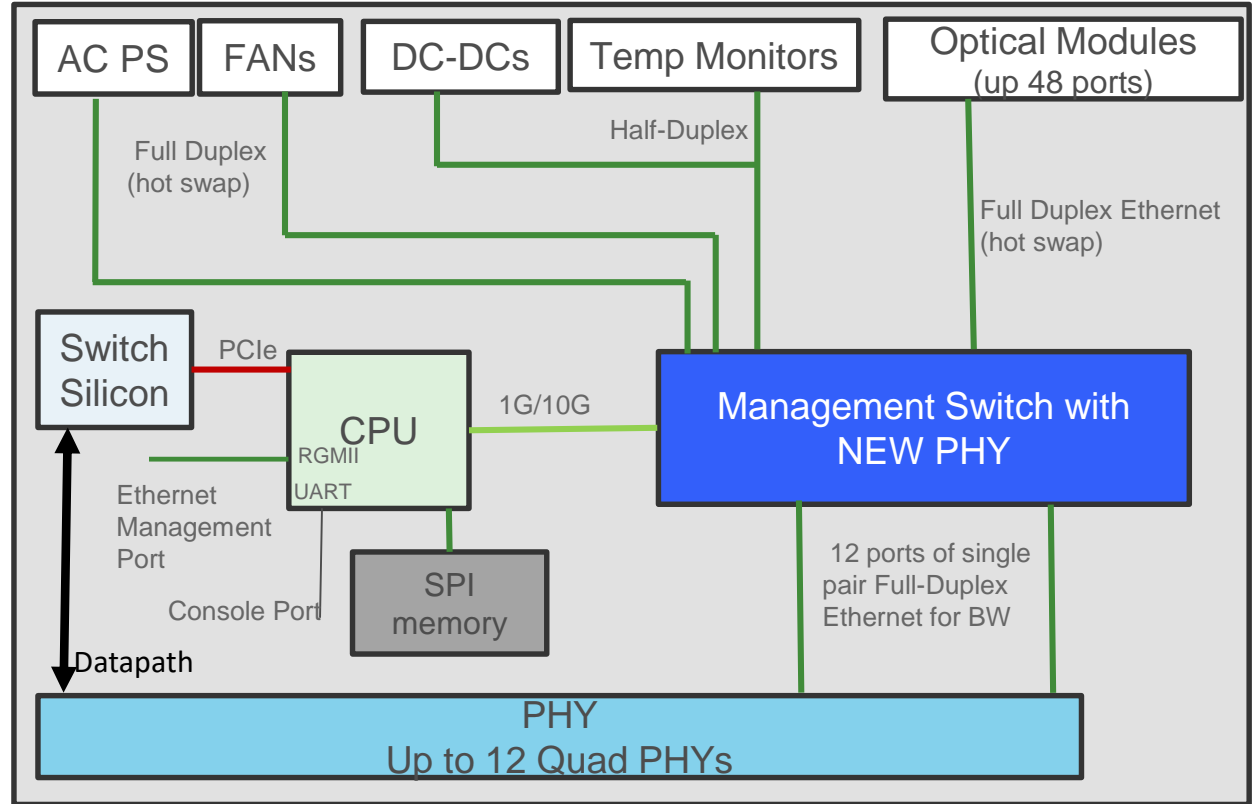
Nov 2017

Why 10BP?

- 2-wire: Reduce inter-connect
- 10Mb/s is higher usable bandwidth than we have today
 - 100Mb/s may be interesting in future
- Ethernet based – one common driver
- Switched Point-to-Point
 - for BW or hot-puggable devices
- Half-Duplex for low bandwidth and fixed devices
 - e.g., DC-DC, Temp monitors etc.

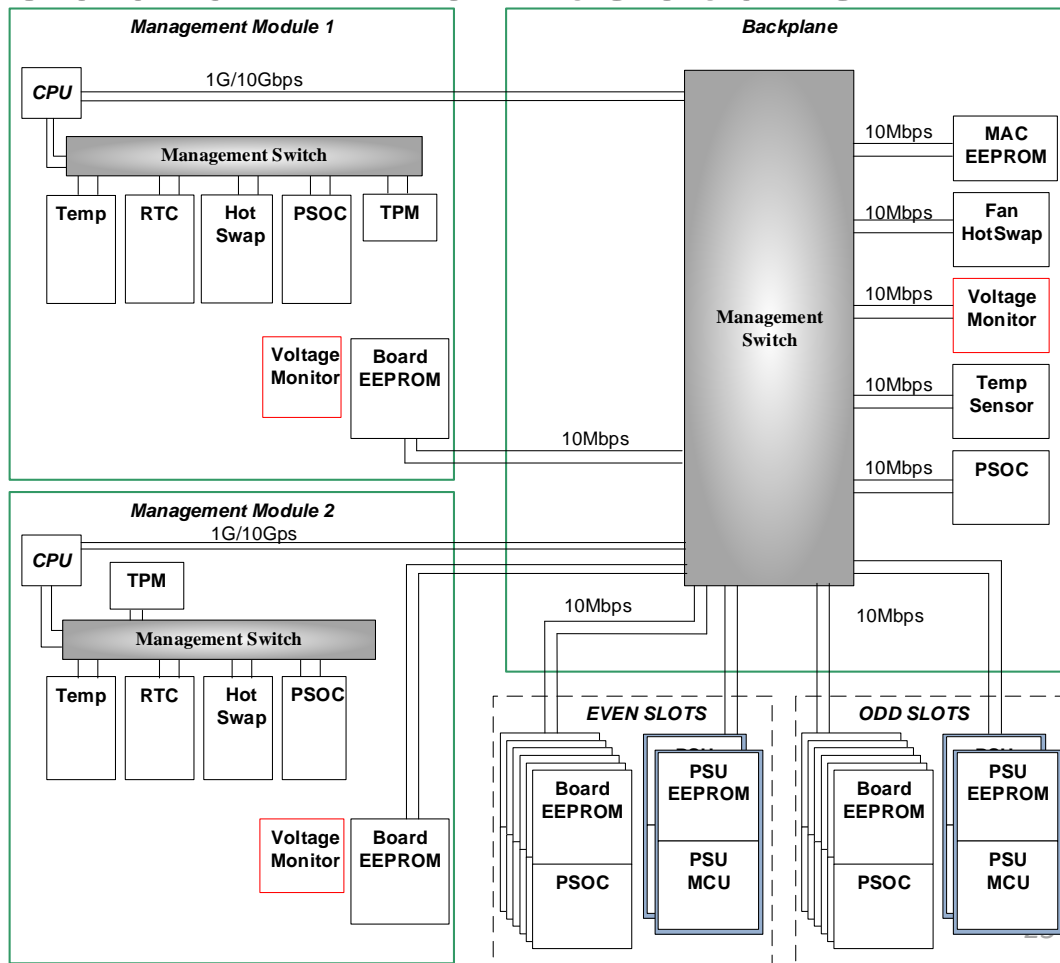
Desired Fixed Architecture

- Number of different I/O types reduced by using
- CPU I/O limitation for I2C/MDO removed
- Point-Point Ethernet for BW and or Hot-pluggable interfaces
- Ethernet Driver



Desired Modular Architecture

- RBOM Reduction
- Point to Point Connectivity
- Single SW Driver
- Full Bandwidth Support



Slide: David Tremblay, HPE

Nov 2017


10BP Technology Feasibility

Mandeep Chadha, Microsemi

Lengths/link segment

- IEEE 802.3cg has a short reach link segment adopted
 - Based on cabling and connector measurements
- Consistent with needs of backplane channels

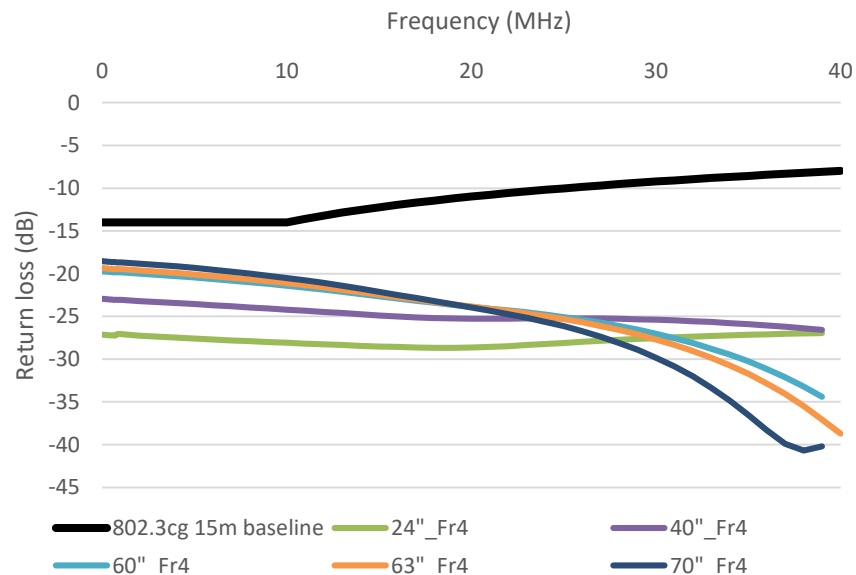
IEEE P802.3cg 15m Link Segment




- Adopt the equations on slide 18 of http://www.ieee802.org/3/cg/public/Sept2017/DiBiaoBergner_01c_0917.pdf as a baseline for the 10SPE short reach link segment.

IL <	$1 + 1.6 (f-1)/9$ dB	$f=0.3 \dots 10$ MHz
	$2.6 + 2.3 (f-10)/23$ dB	$f=10 \dots 33$ MHz
	$4.9 + 2.3 (f-33)/33$ dB	$f=33 \dots 40$ MHz
RL >	14 dB	$f=0.3 \dots 10$ MHz
	$14 - 10 \cdot \text{LOG}_{10}(f/10)$ dB	$f=10 \dots 40$ MHz
MC >	30 dB	$f=0.3 \dots 20$ MHz
	$30 - 20 \cdot \text{LOG}_{10}(f/20)$ dB	$f=20 \dots 200$ MHz

Backplane channel characteristics



No issues with meeting proposed 10SPE
Return Loss specifications



Insertion Loss spec proposed for 10SPE
can be met for channels up to 40" FR4

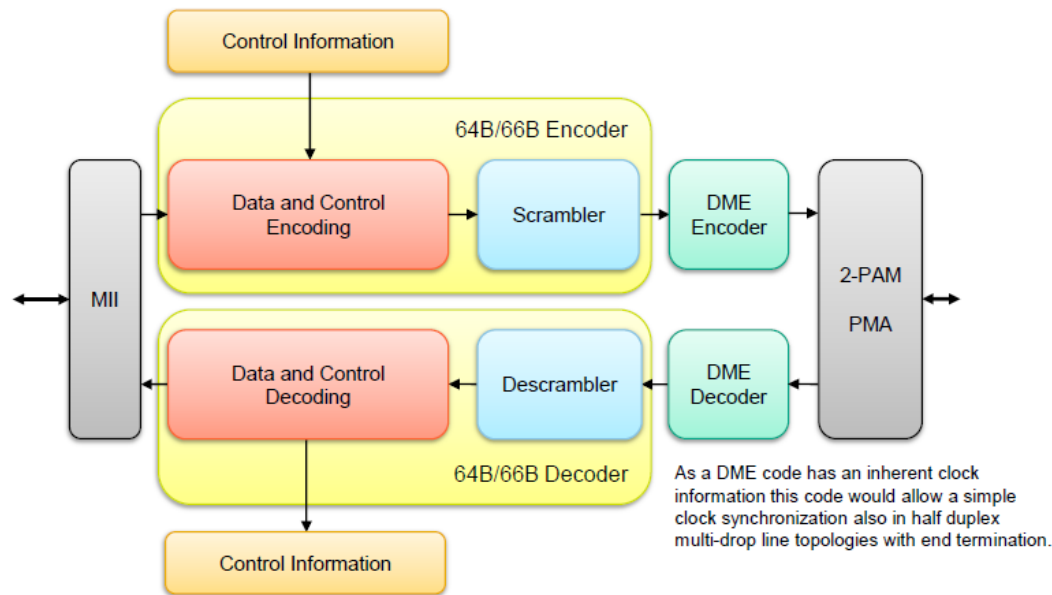
10Mb/s Ethernet in Micros/FPGAs

- Low insertion loss of channel gives lots of choices
- Many encoding types available
 - Differential Manchester, NRZ, etc.
 - All are low-complexity and have been implemented in FPGA format
- Duplexing method is new for micros/FPGAs
 - E.g., echo-cancelled full-duplex or half-duplex on medium
 - Already under study in 802.3cg for short-reach channel

10BP Complexity Estimate

- Proposals under consideration for PHY implementations of short-reach 10SPE use easily realizable analog and digital circuits using standard logic processes
- PMA
 - Differential line driver, analog hybrid/EC, clock sync
- Digital
 - TX – PCS, Scrambler, Encoder
 - RX – Decoder, Descrambler, PCS
- Construct as a macro that can be instantiated in ASIC or FPGA

10SPE short reach PHY proposal in slide 16 of Graber_3cg_12_0717.pdf



10BP - Why Now?

George Zimmerman, CME

Why Now?

- Interest from Network Equipment and Computer OEMs has created the potential for large volume short-reach interconnects
- Leverage investment in standardization of 10Mbps single-pair technology
 - Relevant PHY experts are gathered in IEEE 802.3cg
 - While 10Mbps Ethernet systems are old hat, there aren't any single-pair versions standardized
 - Open and common interoperable specifications simplify the market
 - Requirements and needs are consistent with 802.3cg short-reach objective
- Every time I turn around, there's a new application being proposed
- BUT: Existing 802.3cg project documentation is specific to “single balanced twisted pair copper cabling”

The Rub: 802.3cg PAR Scope, very specific

5.2.b. Scope of the project: Specify additions to and appropriate modifications of IEEE Std 802.3 to add 10 Mb/s Physical Layer (PHY) specifications and management parameters for operation, and associated optional provision of power, on single balanced twisted-pair copper cabling.

- Limitations: PHY & optional powering project, Rate = 10Mbps
- The medium is single balanced twisted-pair copper cabling
 - 2-pair, 4-pair, PCB backplane, parallel pairs, single ended are out of scope
- The Physics: addressing backplane applications is a natural
 - Including PCB “pairs”: The electrons don’t care about twisting the wires or read PAR documents...
- A standard can fill a broader need with 802.3cg solutions, avoiding proprietary extensions

Process

- Call for Interest on 10Mbps Backplane Ethernet to be held in November
- Study group will meet to propose PAR and CSD modifications for 802.3cg documents
 - Expect most work will be done quickly, by close of January meeting
 - Study group can pre-discuss any needed objectives modifications for 802.3cg
 - 802.3 approval of modified P802.3cg project documentation, targeted by March 2018
- Meanwhile, 802.3cg continues its work on short-reach and long reach PHYs
 - Contributors work offline to prepare and build consensus for any necessary text
 - Time to fold in anything necessary for backplane before May WG ballot milestone
- P802.3cg motion showed unanimous support: [P802.3cg Sept'17 motion #15](#)
- We've seen this before - adding 25GBASE-T to 40GBASE-T IEEE P802.3bq
 - Smaller scope of change than in 802.3bq – no new PHYs expected, no speed change (MAC interface, registers, or frequency translations)

Q&A

Presenters

Amrik Bains

George Zimmerman

Jon Lewis

Mandeep Chadha

Expert Panel

David Tremblay

Cisco

CME Consulting

Dell EMC

Microsemi

HPE

Straw Polls

Call-for-Interest Consensus

- Should a study group be formed for “10Mb/s Backplane Ethernet”?
- Y: N: A:
- Room count:

Participation

- I would participate in a “10Mb/s Backplane Ethernet” study group in IEEE 802.3
 - Tally:
- My company would support participation in a “10Mb/s Backplane Ethernet” study group
 - Tally:

Future Work

- Ask 802.3 at Thursday's closing meeting to form study group
- If approved:
 - 802 EC votes on Friday to approve the formation of the study group
 - First study group meeting would be during the January 2018 802.3 interim meeting (in Geneva)

End