

How many connections for 10GBASE-LRM?

Piers Dawe

Agilent

Lew Aronson, Jim McVey

Finisar

Sudeep Bhoja

Big Bear

John Ewen

JDSU

Martin Lobel

Intel

Petre Popescu

Quake

Abhijit Shanbhag

Scintera

Nick Weiner, Ben Willcocks

Phyworks

Statement of problem

- Objective creep and over-engineering
 - Danger of permanent cost and heat burdens
 - Difficult to get to high port density XFP based equipment
- Need to specify to the mainstream, not the corner cases

Over-engineering more problematical than previous optical Ethernet

- Our optical channel is hugely varied and we are squeezing a lot out of it
- No cheap margin as in Gigabit
- Increasing port densities
- Compare 10GBASE-E
 - The introduction to long haul, not the ultimate long haul PMD
 - Was deliberately not over-engineered (e.g. 80 km would have been nice, but...)
- **Need to stop chasing the tail of a distribution**

Issues in more detail

1. TP3 stressors in D1.1 (three emulated channels) do not match modelling information
 - Based on Gen54YY, not current Gen67YY
 - Not based on the 2-launch strategy agreed at last meeting
 - See ewen_1_0503.pdf
2. Stressor search methodology should address poor correlation between PIE-D and real equaliser penalty
 - See ewen_1_0503.pdf
3. **Topology of link (input to modelling) needs review**
 - **Discussed on following slides**


Connectors matter more

- Mechanical tolerances give lateral offset across connection
 - See Al Brunsting's presentation emailed by Jonathan King
- Connectors near the transmitter can degrade the channel's frequency performance
- Multiple connectors would make this worse
- Connectors near the receiver might cause modal noise but budget allows for this
- Connectors are more of an issue at 10G than at 1G
 - (but connectors in 1G not as deeply studied as in 802.3aq)

Summary of connector effects


Effect	In MMF	In SMF	Comments
Loss	Yes	Yes	But not the issue here
Reflection	Yes	Yes	Less serious if multimode laser or fiber
Modal noise	Yes	No	
Mode coupling	Yes	No	Degrades channel freq. response – may be offset by diversity of two fibers

Channel topology




- Supported topology per standard is point-to-point
- More than 90% of links are like this
 - Don't know how much more than 90%
- Note, just one intermediate connector before the long cable

Channel topology 2




- Supported topology per standard is point-to-point.
>90% of links are like this
- End-to-end length ≤ 300 m
- 10G allows more connectors if same total performance e.g overall connector loss
 - this might be an oversight for MMF
- 1G allows more connectors if same total e.g connector loss
- Connectors between transmitter and long run are important
 - Connectors near receiver hardly matter

Alternative topology



- This topology is supported by ANSI/TIA/EIA-568-B.1-2001, Commercial Building Telecommunications Cabling Standard
- Less likely for 10G use as 10G doesn't go to desktop
- Supported with restrictions by Optical Ethernet
- If two similar runs, fiber diversity helps considerably
- With good fibre, shorter links and/or better connectors it will work often enough

“Modal noise test” topology



- Note, two intermediate connectors before the long cable
- This topology designed to be bad, and to create enough modal noise to be observed
- Not a service channel
- Not compliant (unless connector offsets are low)

Use scenarios

A. Campus backbone

- Not likely, 300 m generally not enough

B. Building backbone

- Vast majority with just 2 intermediate connectors

C. Horizontal (to desktop)

- Not likely, neither 10G nor fiber go to desks

D. Combined backbone & horizontal, passively connected

- TIA/EIA 568 B and ISO/IEC 11801's "Centralised Optical Architecture"
- Not widely used. 3 intermediate connectors, could be more

E. Data centers

- Up to ~50m, 4 intermediate connectors
- Committee has been addressing B above
- Can also offer useful performance on E


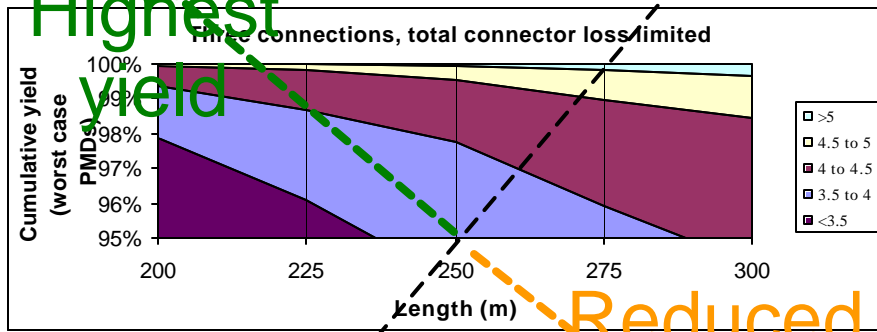
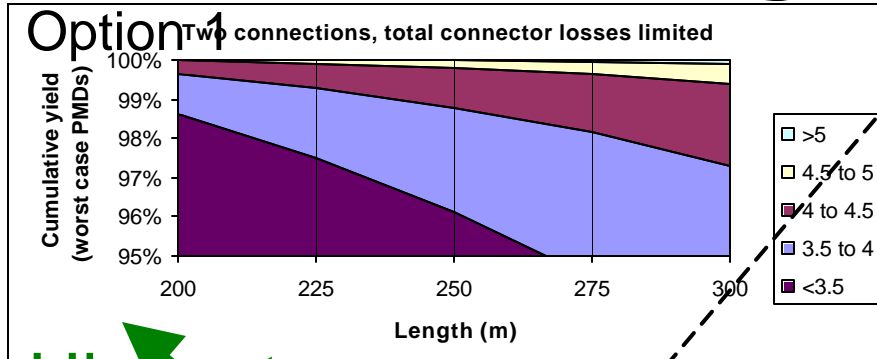
Green – relevant

Red – not relevant

Options for our connector scenario

1. Just two intermediate connections as shown in fig 38-7 or 52-14 (not counting PMD connectors at DTE face plate)
 - Only one before long run (main cable).
 - Might be unhelpful to a minority of users
2. Allow up to three intermediate connections but within overall connector loss spec
 - As D1.1. Generous to a minority of users, imposes unnecessary thermal/performance burden on the majority
3. With an extra connector before long run
 - Formally not correct, similar to option 2 –onerous
4. (Data center – short, more connectors)
5. Design to option 1, give guidance for options 2 perhaps with reduced reach or tighter connector specs and option 4

Yields vs. length and connectors



Monte Carlo simulations. ~parabolic curves. See backup for another view of data. Loss limit makes no difference for 2 connectors.

Effect of the third connector is similar to 50 m change of length
Each dB costs more than the one before

Middle chart aligns to D1.1

Over-engineered: very few channels are a full 300 m with three end-of-life connectors

Propose 802.3aq specs to top-row scenario

Recommendation for 802.3aq

5. Design to option 1 (just two intermediate connections),
give guidance for option 2 (up to three intermediate connections but within overall connector loss spec)
 - perhaps with reduced reach or tighter connector specsand give guidance for option 4 (short, several connectors)

Backup slides

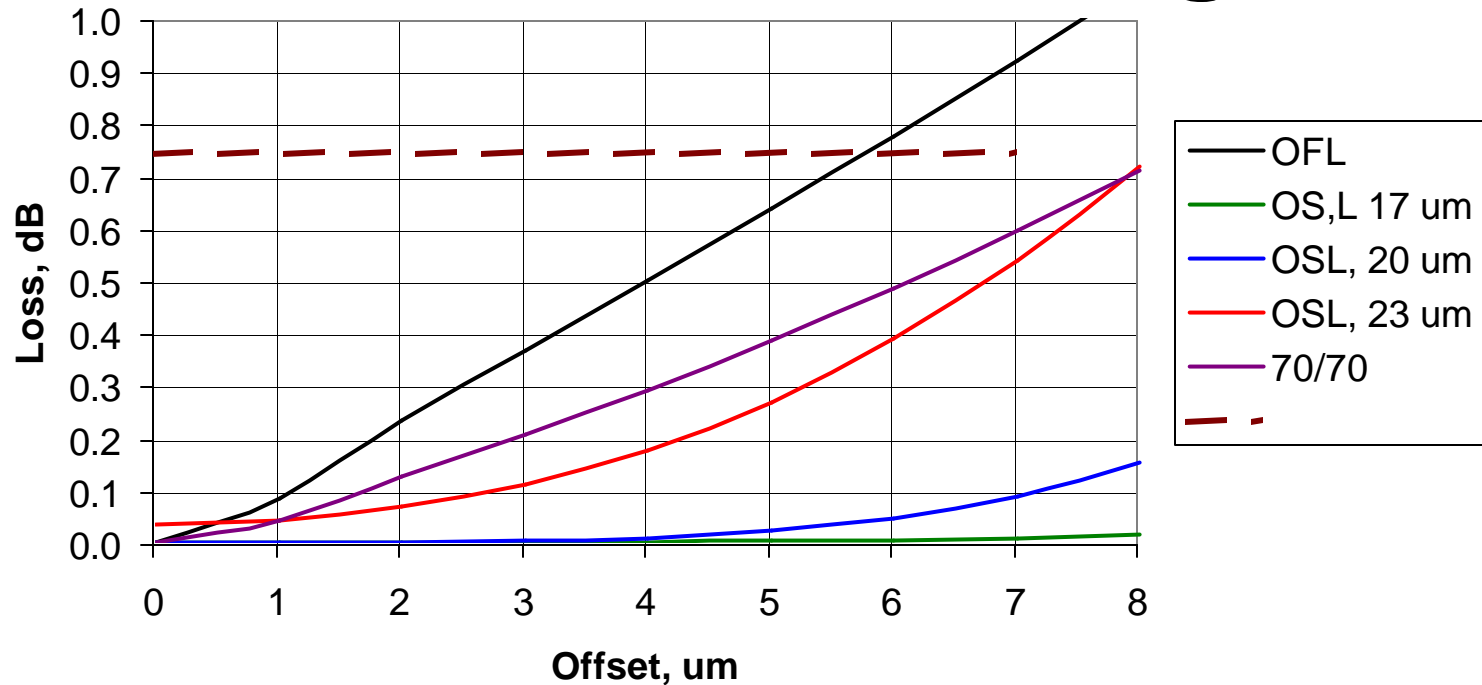
What determines connector quality?

- Widely accepted model per Pepeljugoski: Rayleigh distribution, mean 3.58 μm , 95% < 7 μm
- Brunsting showed medians 2.5 to 3 μm , 95% < 6 μm
- Connector quality has gradually improved over the years
- There are different grades of connector
- Might reconsider if 7 μm is too pessimistic
 - If one can afford 10G ports, one can afford new/good fiber terminations!
 - Question is logistics, not cost

How is connector loss defined?

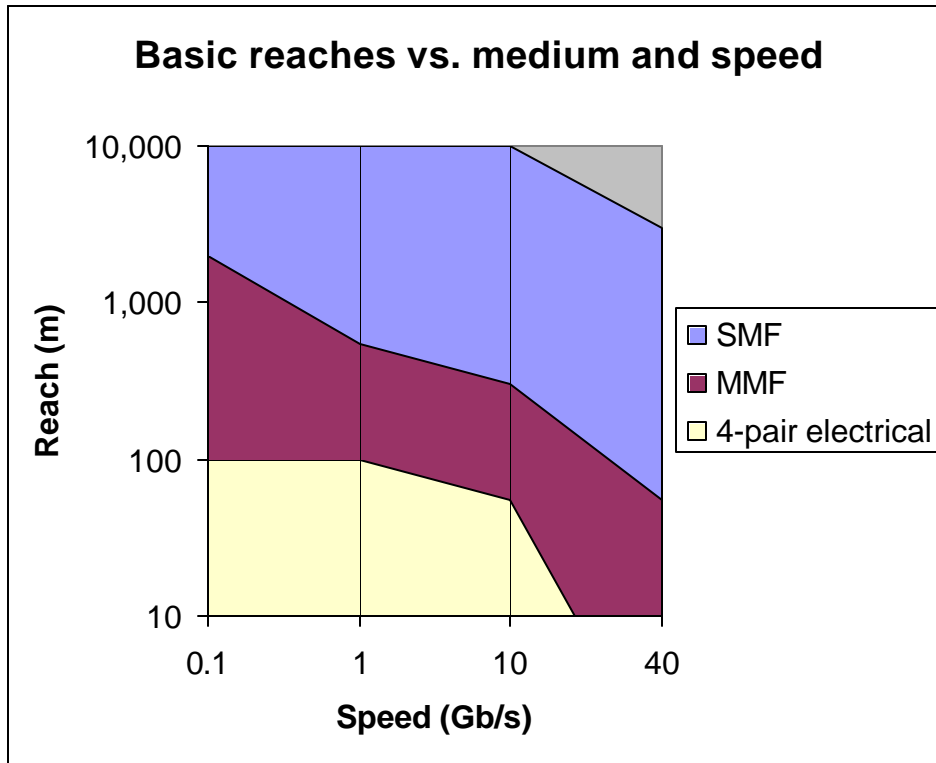
- Loss is divided into two parts:
 - Loss due to offset
 - Light spills out of the side
 - “Parameter mismatch”
 - S21 across connector < 1 due to impedance mismatch of two glass waveguides
 - Rule of thumb: 2/3 is due to offset
- Mode-power distribution for loss spec is defined
 - Since 1995, OFL mode power distribution or 23 μm OSL – significant power near outside of core
 - Before 1995, equilibrium mode-power distribution or 70/70 distribution – rather less power near outside of core, therefore allowed worse offset
 - Which should we use, in 2005?
- Connector loss is NOT defined for the particular launches we will use
 - Actual loss with these launches is $<$ spec loss

Connector loss vs. offset for different mode-fillings




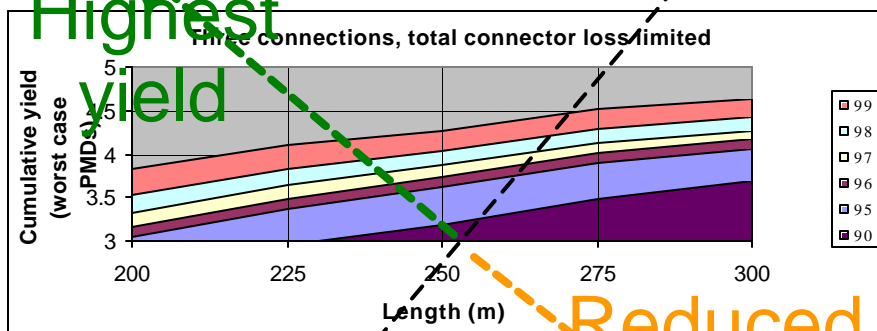
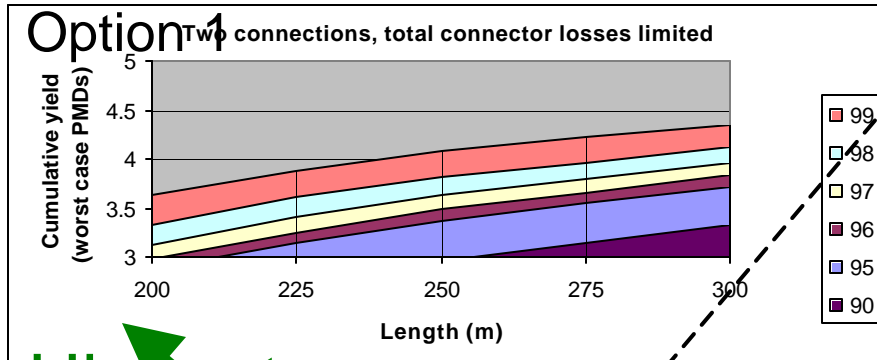
- For connector offsets < 7 um, actual 10GBASE-LRM loss in service (red, blue, green) $<$ measured loss with loss meter (black or magenta)
- Parameter mismatch ignored here

Reach comes down as data rate goes up



- Every time the speed goes up, the reach goes down
- For next generation, use SMF
 - Which also will tolerate more connectors

Yields vs. length and connectors – alternative view



Highest yield

Reduced yield

Penalty rises ~linearly with length, rises **rapidly** with %

Effect of the third connector is similar to 50m change of length

Each dB costs more than the one before

Middle chart aligns to D1.1

Over-engineered: very few channels are a full 300 m with three end-of-life connectors

Propose 802.3aq specs to top-row scenario

References

- Al Brunsting & Rick Pimpinella, “Lateral offsets for multimode fiber (MMF) connectors”
<http://ieee802.org/3/10GMMFSG/email/msg00568.html> and
<http://ieee802.org/3/10GMMFSG/email/msg00569.html>
- Petar Pepeljugoski et al., “Development of System Specification for Laser-Optimized 50- μ m Multimode Fiber for Multigigabit Short-Wavelength LANs”, IEEE J. Lightwave Technology **21**(5), pp 1256-1275, May 2003