Scaling Optical Networking Capacity: Options and Solutions

Peter J. Winzer
Nokia Bell Labs, Holmdel, NJ

NOKIA

A Big *Thank You* for Lots of Discussions and Input to These Slides

Heidi Adams
S. Chandrasekhar
Xi (Vivian) Chen
Junho Cho
Andy Chraplyvy
Ronen Dar
Randy Eisenach
Nick Fontaine
Kyle Guan
Steve Grubb
Andrew Lord
David Neilson
Greg Raybon
Roland Ryf
Bob Tkach
Szilard Zsigmond

You Can Find More on the Topics of This Talk Here:

JOURNAL OF LIGHTWAVE TECHNOLOGY, VOL. 35, NO. 5, MARCH 1, 2017 1099

From Scaling Disparities to Integrated Parallelism: A Decathlon for a Decade

Peter J. Winzer, *Fellow, IEEE*, and David T. Neilson, *Fellow, IEEE*

(Invited Tutorial)

Fiber-Optic Transmission and Networking: The Previous 20 and the Next 20 Years

PETER J. WINZER*, DAVID T. NEILSON, AND ANDREW R. CHRAPLYVY

Nokia Bell Labs, 791 Holmdel Road, Holmdel, NJ 07733, USA
peter.winzer@nokia-bell-labs.com

NOKIA

Network Traffic is Growing Tremendously

If the Past Informs the Future: Expect Exponential Network Traffic Growth

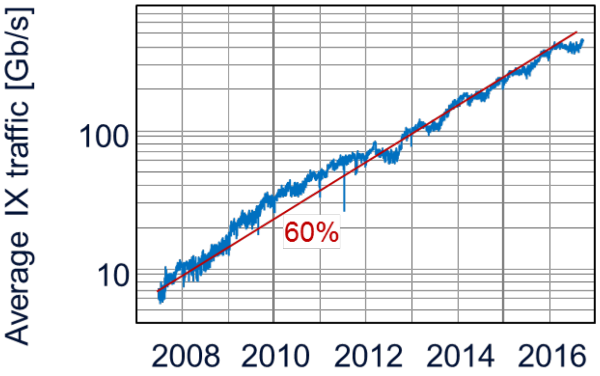

(b) Seattle Exchange (SIX) [20]

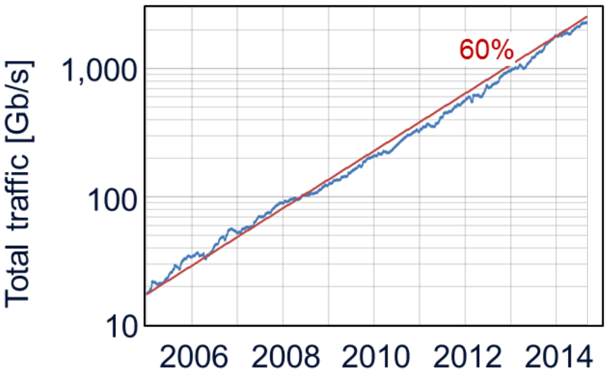

(c) Total broadband demand (UK ISP) [3]

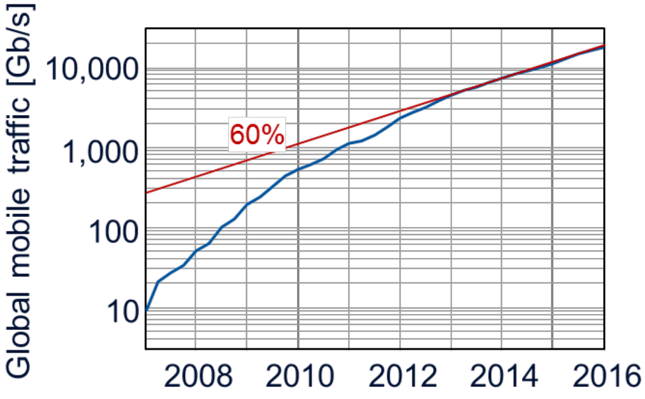

(e) Global mobile data traffic [22, 23]

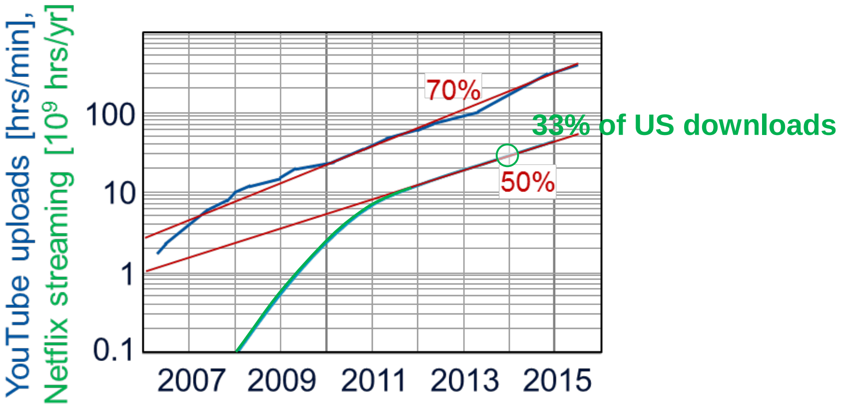

(f) Video traffic (YouTube, Netflix) [24, 25]

NOKIA

Looked at From a Global WDM Transponder Sales Point of View

➔ Global Network Traffic is Growing Around 45%

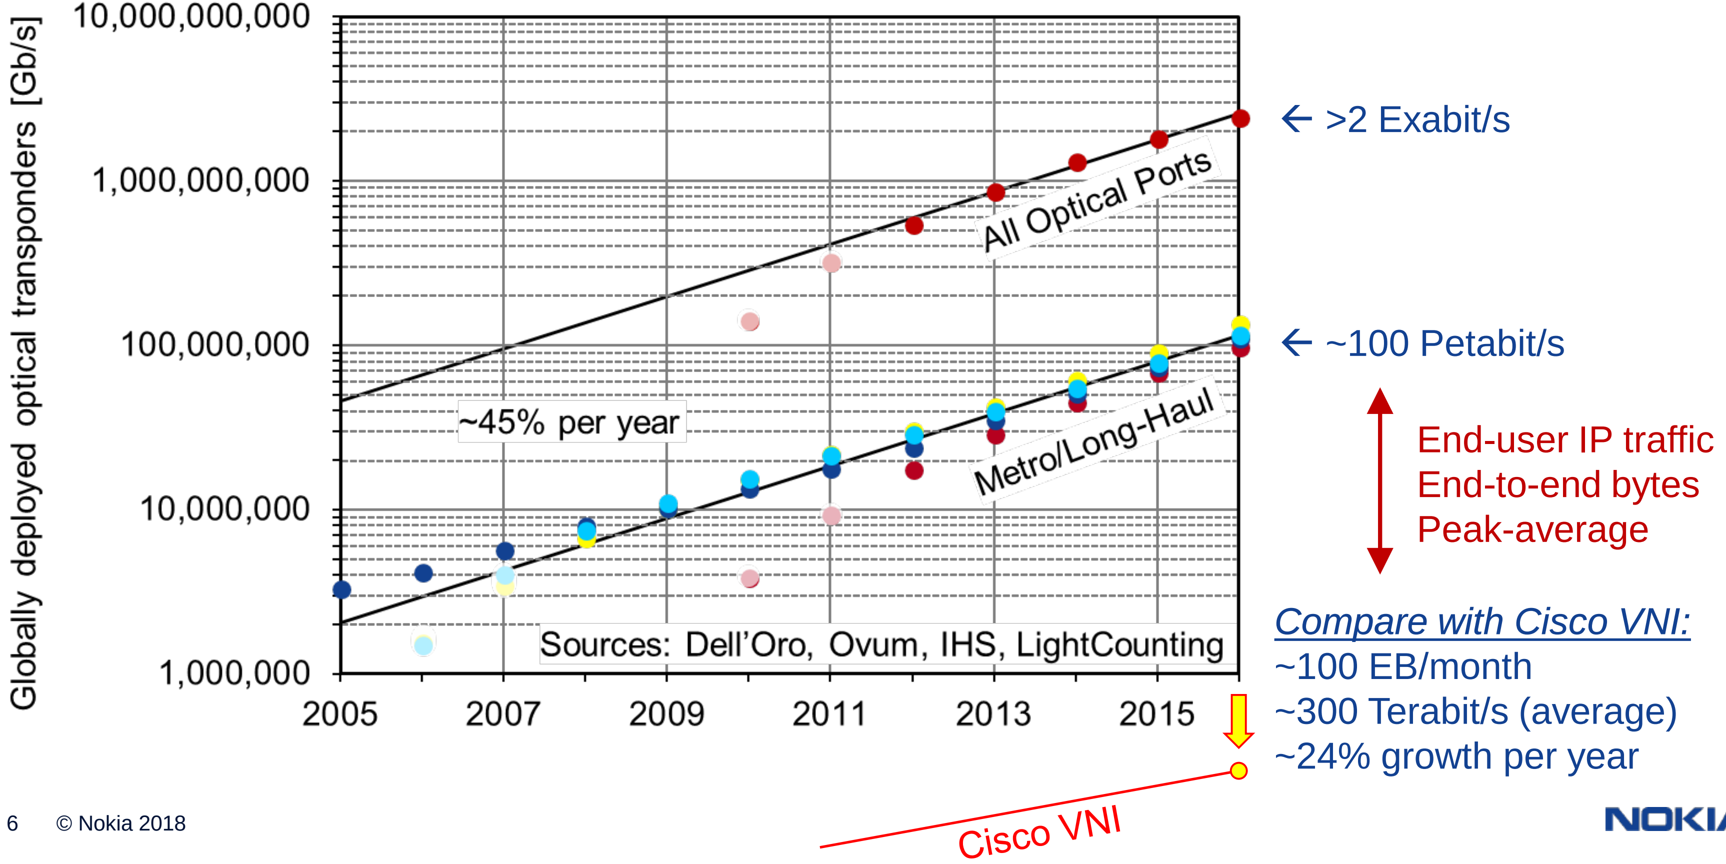

Compare with Cisco VNI:
~100 EB/month
~300 Terabit/s (average)
~24% growth per year

NOKIA

Traffic Growth Driven by Compute, Storage, and Access Technologies

Information Generation, Consumption, Processing

Technology scaling	Exponential trend period	CAGR
Supercomputer performance	1995 – 2017	90%
Microprocessor performance	1980 – 2017	40% - 70%
Storage capacity	1980 – 2017	60%
Core router capacity	1985 – 2017	45%
Wireless interfaces	1995 – 2017	60%
Fixed access interfaces	1983 – 2017	40 - 55%

~60%

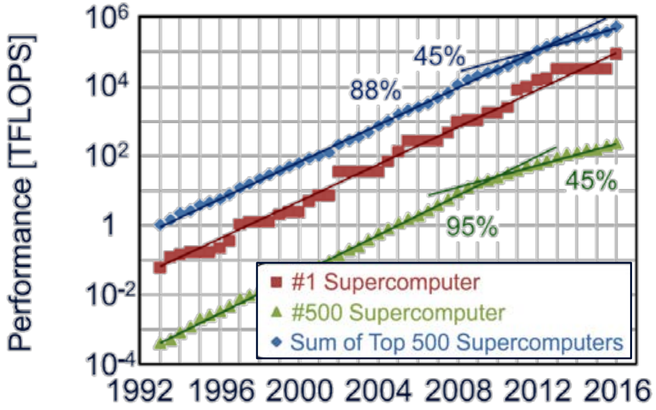

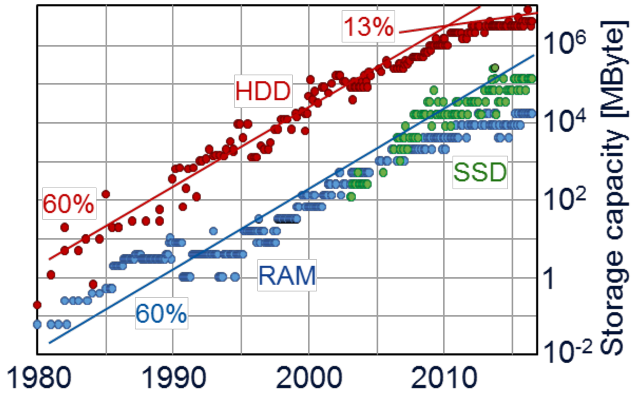

NOKIA

Optical Networking: Interface Rates

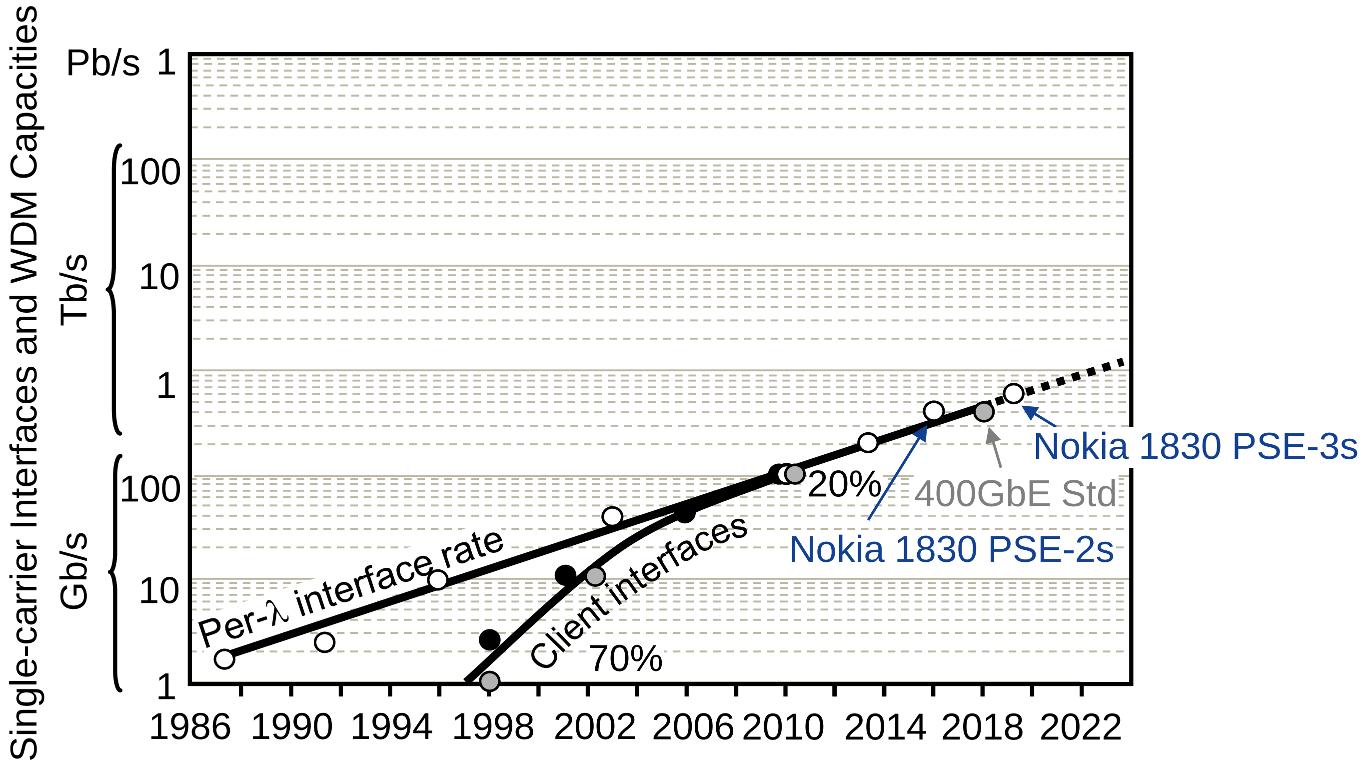

NOKIA

Optical Networking: Interface Rates: ~10T Client Interfaces by ~2025?

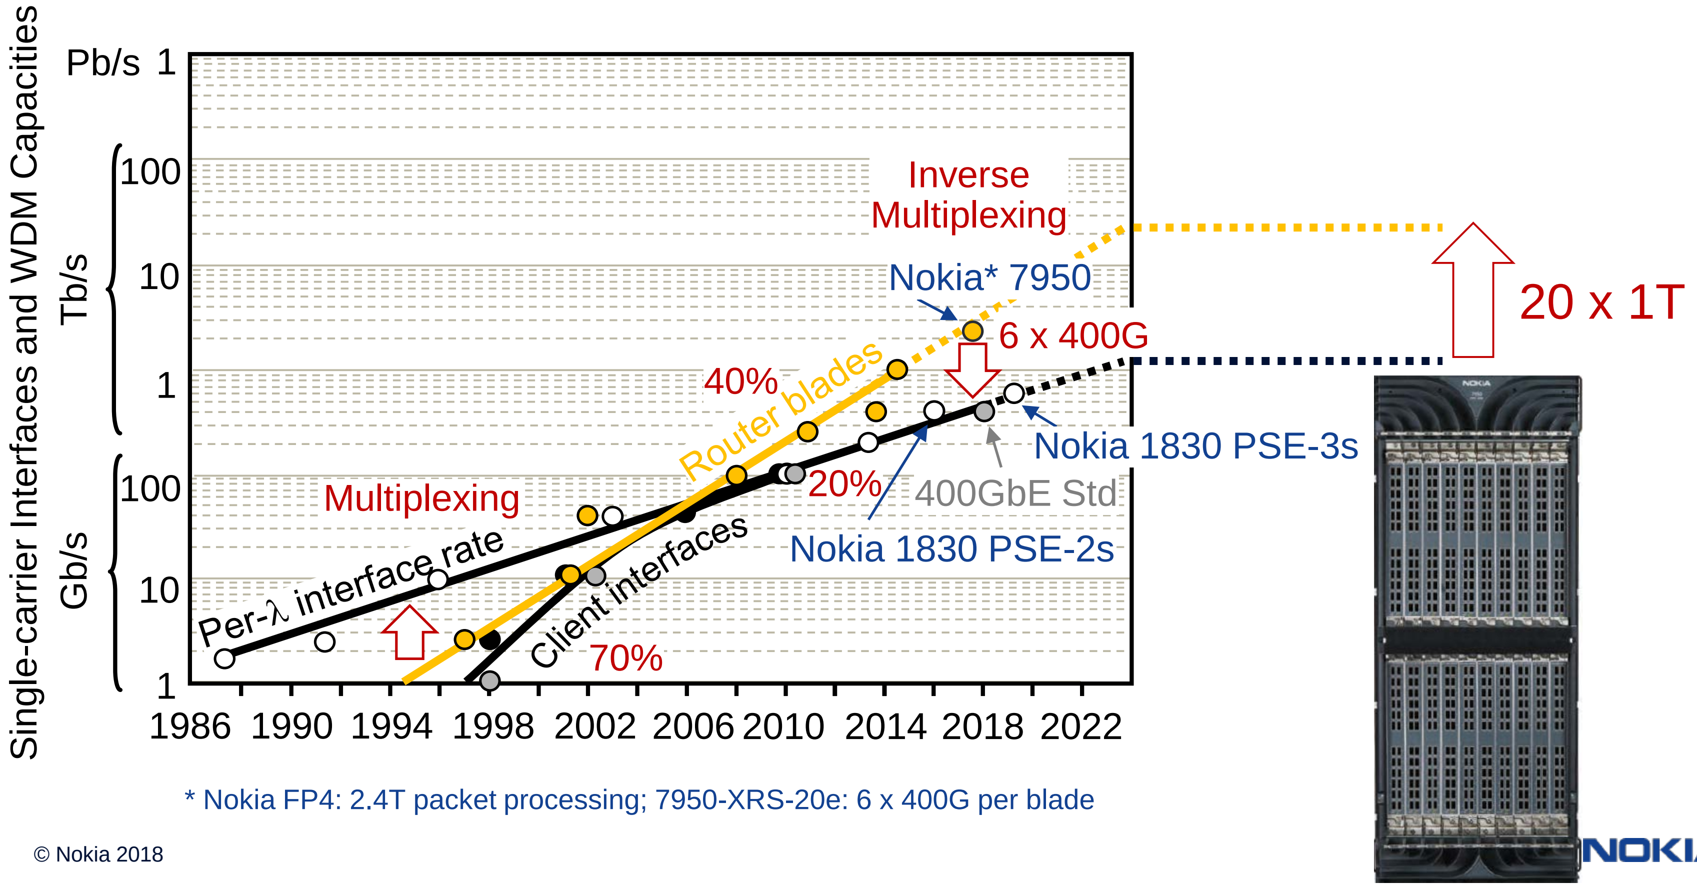

* Nokia FP4: 2.4T packet processing; 7950-XRS-20e: 6 x 400G per blade

NOKIA

The Exact Same Scaling Disparities are Found in Switch Chips

Per-Lane Speed vs. Switching Capacity

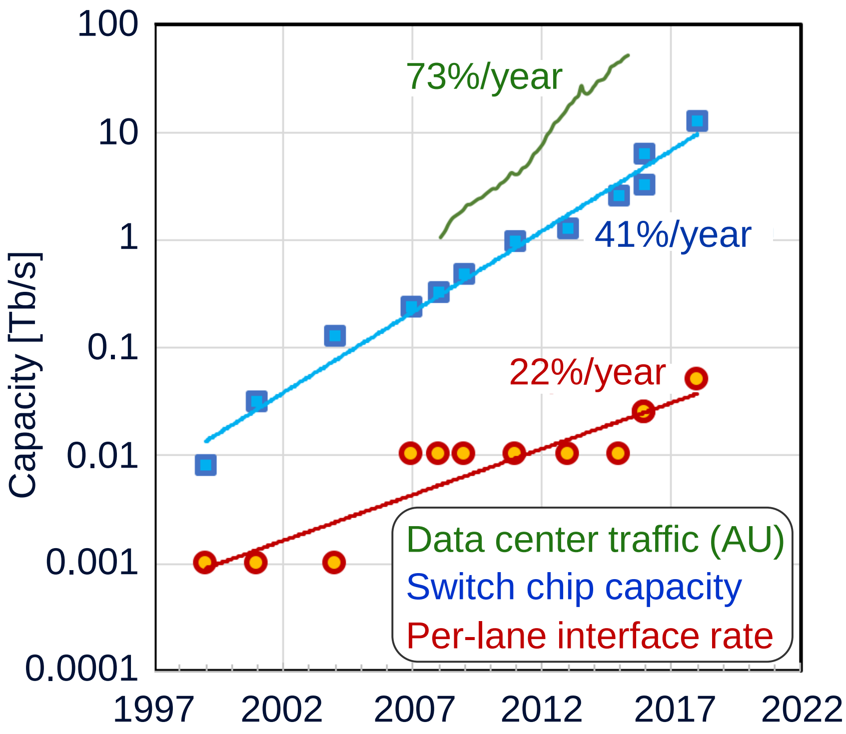

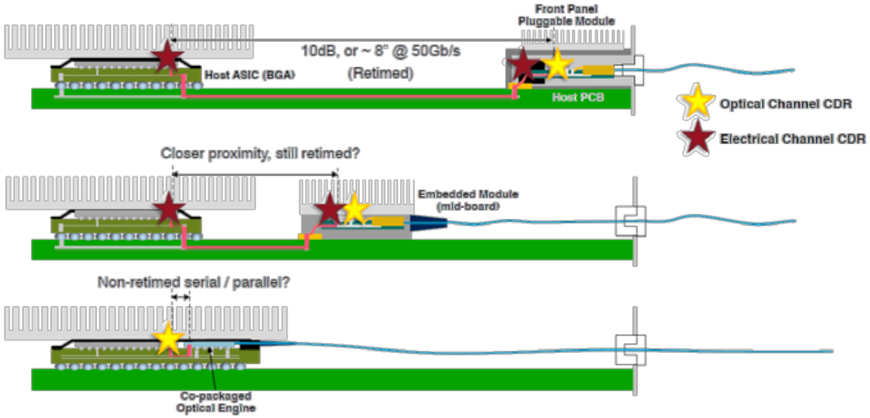

Figure after:
[R. J. Stone, OFC'17, Th3G.5] - Broadcom
[A. Singh et al., Sigcomm'15, 183] - Google

Wavelength Division Multiplexing: Petabit/s systems by ~2025?

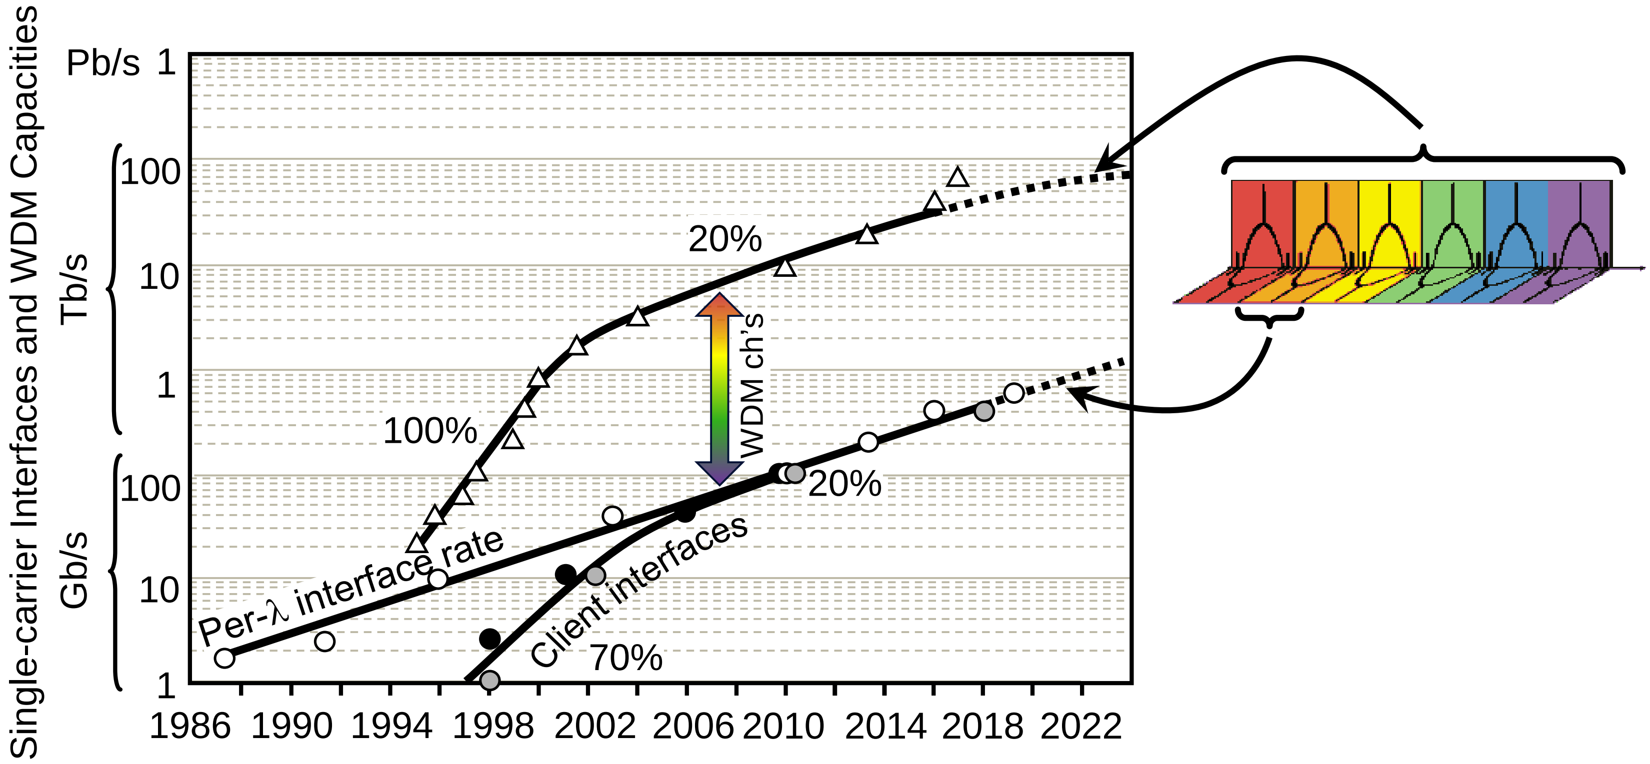

NOKIA

Growing Disparity Between Generation and Transport of Information

Technology scaling	Exponential trend period	CAGR
Supercomputer performance	1995 – 2017	90%
Microprocessor performance	1980 – 2017	40% - 70%
Storage capacity	1980 – 2017	60%
Core router capacity	1985 – 2017	45%
Switch chip capacity	1998 – 2018	40%
Wireless interfaces	1995 – 2017	60%
Fixed access interfaces	1983 – 2017	40 - 55%
Router interface speed	1980 – 2005	70%
	2005 – 2017	**20%**
Transport interface speed	**1985 – 2017**	**20%**
Per-lane chip interface speed	**1998 - 2018**	**20%**
WDM capacity per fiber	1995 – 2000	100%
	2000 – 2017	**20%**

Information Generation, Consumption, Processing: 40%-90% (~60%)

Information Transport: ~20%

!

5 years: 4x disparity

10 years: 17x disparity

NOKIA

How did we get to where we are ? Where do we need to go long-term ?
What are the limits ? How do we get to where we need to be ?

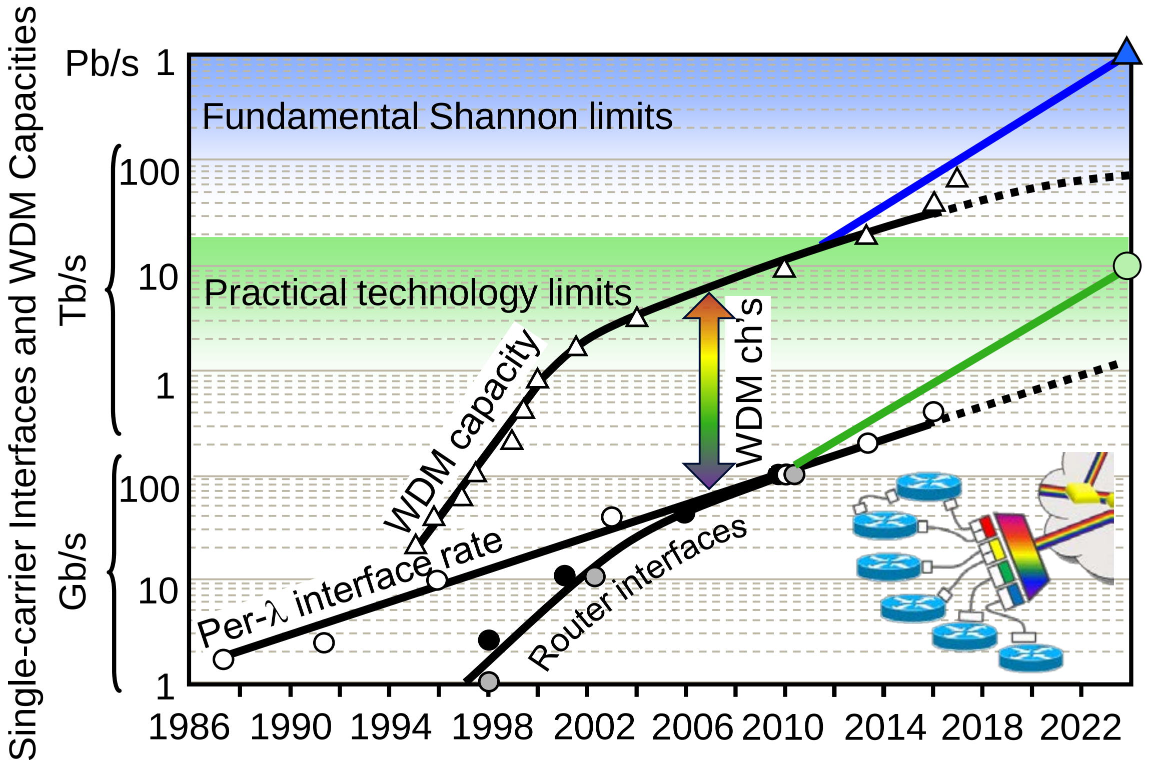

1-Petabit/s systems by 2024

"only" at 40% growth

10-Terabit/s interfaces by 2024

NOKIA

Capacity scaling: What are the options ?

$$C = 2\,M\,B\,log_2(1 + SNR)$$

Polarization | Bandwidth

Spatial paths

Pre-log (multiplexing) factors

Logarithmic (modulation) capacity

NOKIA

Capacity scaling: What are the options ?

$$C = 2\,M\,B\,log_2(1 + SNR)$$

Polarization | Bandwidth

Spatial paths

Pre-log (multiplexing) factors

Logarithmic (modulation) capacity

NOKIA

Fundamental Scalability Problems: Shannon Limits to Optical Fiber Capacity

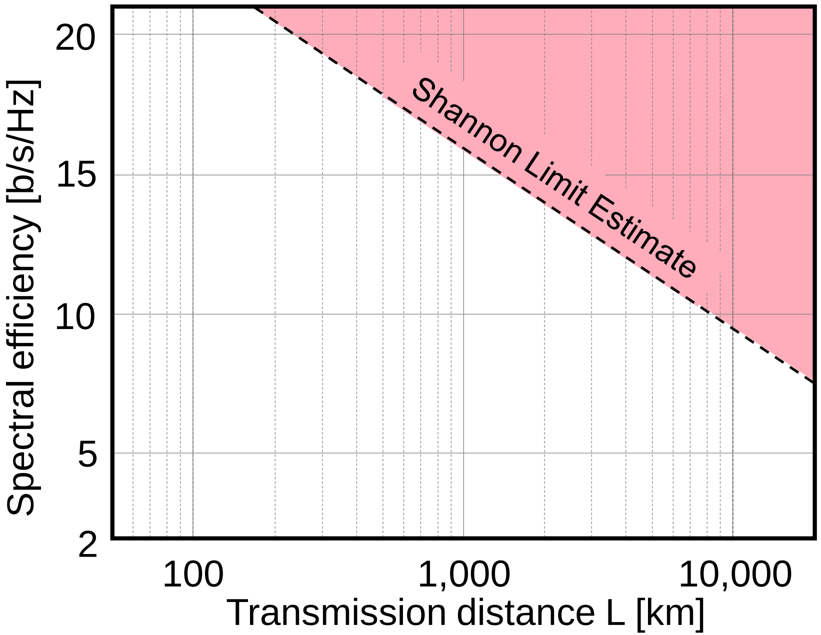

$C \sim \log_2(1+SNR)$

~1/L

Optical amplifier noise

Nonlinear interference noise

[P. Poggiolini et al., J. Lightwave Technol. (2014)]
[R. Dar et al., Opt. Express (2014)]

Record Experiments Approaching Fundamental Limits

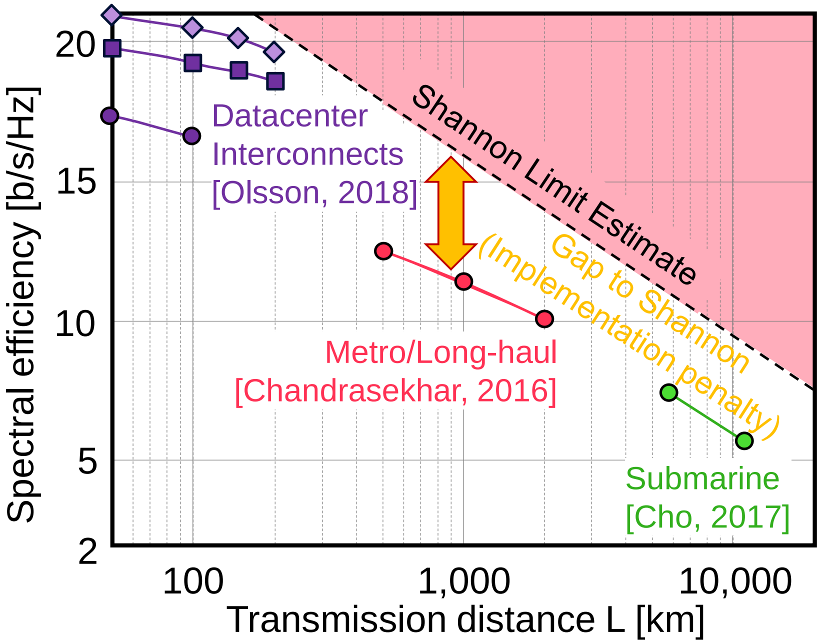

[G. Böcherer et al., IEEE Trans. Commun. (2015)]
[F. Buchali et al., J. Lightwave Technol. (2016)]
[J. Cho et al., J. Lightwave Technol. (2018)]

Probabilistically Shaped QAM

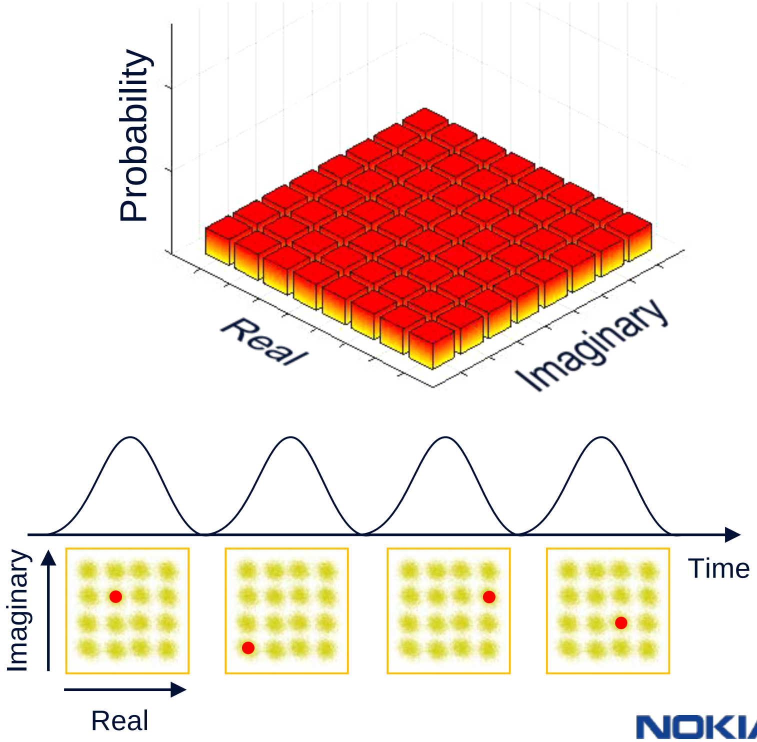

Increasing Capacity by Improving the SNR

$C \sim \log_2(1+SNR)$

- Digital Backpropagation
  (and various computationally simpler approximations)
- Nonlinear Fourier Transform

$$\frac{\partial A}{\partial z} = -\alpha A - j\beta \frac{\partial^2 A}{\partial t^2} + j\gamma |A|^2 A + N$$

Increasing Capacity by Improving the SNR

$C \sim \log_2(1+SNR)$

- Digital Backpropagation
  (and various computationally simpler approximations)

- Nonlinear Fourier Transform

$$\frac{\partial A}{\partial z} = -\alpha A - j\beta \frac{\partial^2 A}{\partial t^2} + j\gamma |A|^2 A + N$$

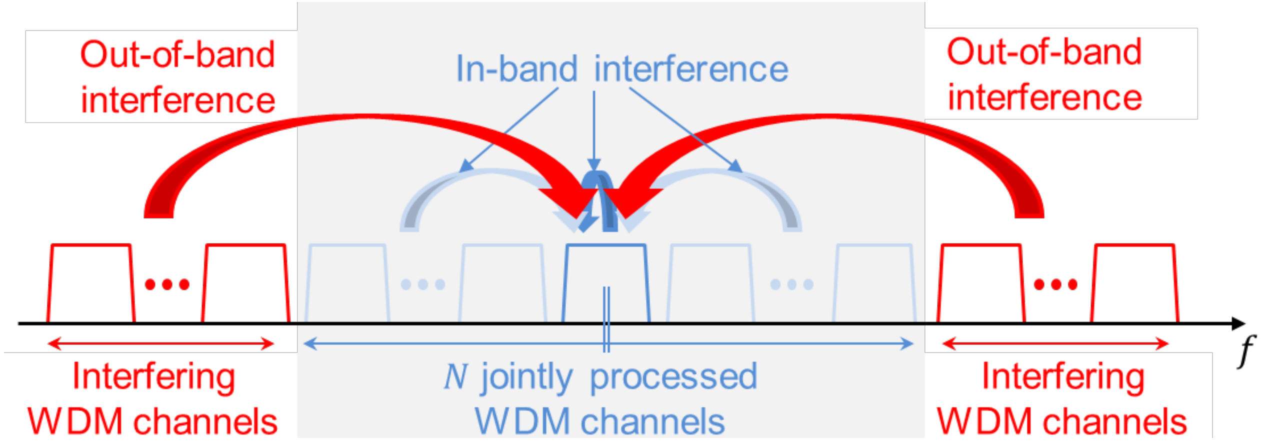

Recent comprehensive reviews:
[Dar and Winzer, J. Lightwave Technol. (2017)]
[Cartledge et al., Optics Express (2017)]

Increasing Capacity by Improving the SNR

$C \sim \log_2(1+SNR)$

- Digital Backpropagation
  (and various computationally simpler approximations)

- Nonlinear Fourier Transform

$$\frac{\partial A}{\partial z} = -\alpha A - j\beta \frac{\partial^2 A}{\partial t^2} + j\gamma |A|^2 A + N$$

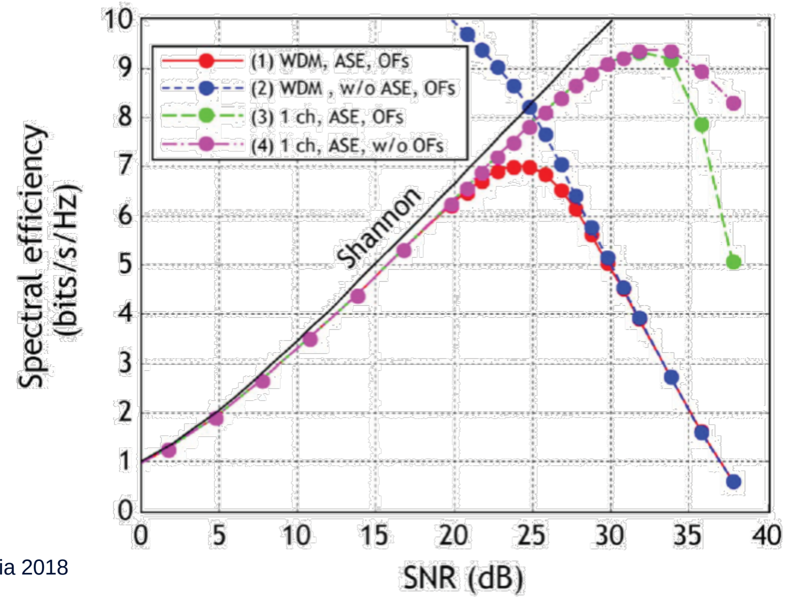

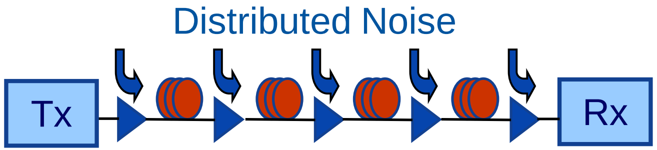

[Essiambre et al., J. Lightwave Technol. (2010)]

NOKIA

Increasing Capacity by Improving the SNR

$C \sim \log_2(1+SNR)$

- Digital Backpropagation
  (and various computationally simpler approximations)
- Nonlinear Fourier Transform

- Low-loss amplification
  (Raman, phase-sensitive)

Example: At 20 dB SNR, what does a 3-dB lower noise figure buy you?

$\log_2(100) \rightarrow \log_2(200)$ … 6.6 b/s/Hz → 7.6 b/s/Hz … 15% more capacity

Don't mess with the SNR !

NOKIA

Increasing Capacity by Improving the SNR

$C \sim \log_2(SNR)$

- Digital Backpropagation
  (and various computationally simpler approximations)

- Nonlinear Fourier Transform

- Low-loss amplification
  (Raman, phase-sensitive)

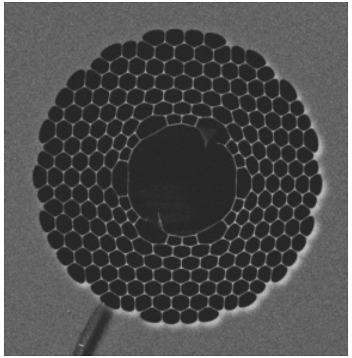

[Petrovich, ECOC 2012]

- Low-loss or low-nonlinearity fiber

  Example: At 20 dB SNR, what does a 10-dB higher launch power buy you?

  $\log_2(100) \rightarrow \log_2(1000)$ … 6.6 b/s/Hz → 10.0 b/s/Hz … 50% more capacity

  - Just in order to _double_ capacity, one needs $\alpha_{dB}\,\gamma$ / 64

Don’t mess with the SNR !

Trading Modulation for Multiplexing

A Good Strategy for a Power-Limited Channel

Where should I put my power?

$$C = 2\,M\,B\,\log_2(1 + SNR)$$

Polarization | Bandwidth

Spatial paths

Pre-log (multiplexing) factors

Logarithmic (modulation) capacity

NOKIA

Trading Modulation for Multiplexing
A Good Strategy for a Power-Limited Channel

Example: At 20 dB SNR, what can I do with 3 dB more *overall* system power ?

$\log_2(100) \rightarrow \log_2(200)$ … 6.6 b/s/Hz → 7.6 b/s/Hz … 15% more capacity

$\log_2(100) \rightarrow 2\log_2(100)$ … 6.6 b/s/Hz → 13.2 b/s/Hz … 100% more capacity

$$C = 2MB \underbrace{\log_2(1 + P/2MBN_0)}$$

Logarithmic Shannon term (modulation)

Linear ("pre-log") term (multiplexing)

Maximum capacity for "infinite multiplexing" in B and/or M:

$$\lim_{M \to \infty} C = (P/N_0)\log_2 e$$

Massive Spatial Parallelism for Cost Efficient Submarine Systems

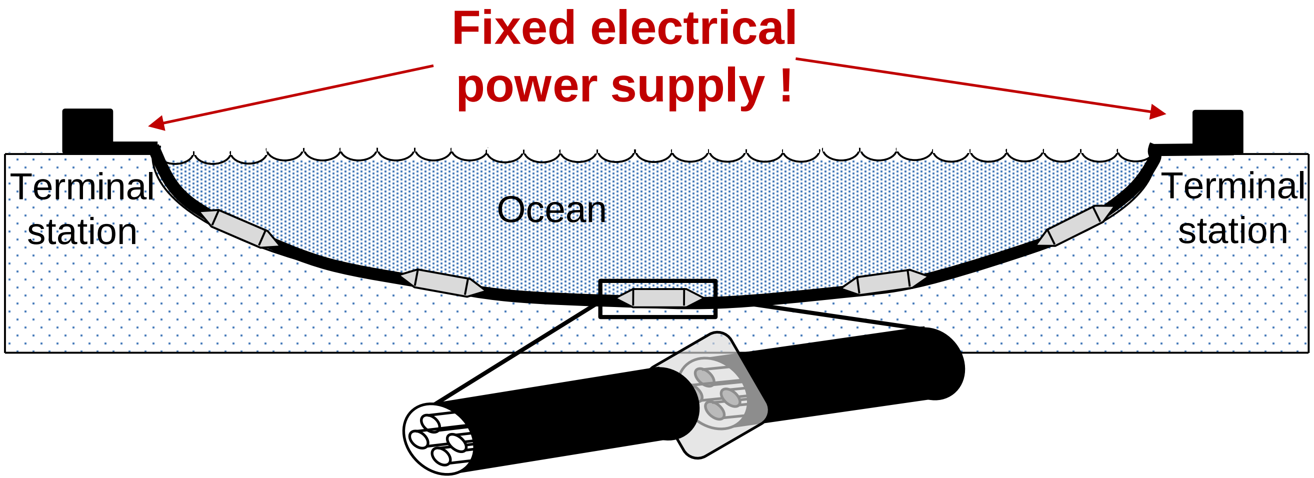

[A Pilipetskii., OFC Tutorial (2015)]
[O. Sinkin, Phot. Technol. Lett. (2017)]
[R. Dar et al., J. Lightwave Technol.(2018)]

Submarine Cable Capacity Model

Power spectral density of the signal

System bandwidth

$$C = 2MB\log_2\left(1 + \frac{PSD}{Nhf\left(Fe^{\frac{\alpha L}{N+1}} - 1\right) + \chi' \log B \, PSD^3}\right)$$

Number of spatial paths

ASE

Nonlinear interference noise (NLIN)

N: Number of amplifiers
L: System length
α: Fiber attenuation
F: Amplifier noise figure
hf: Photon energy
χ': NLIN coefficient

[R. Dar et al., Proc. ECOC, Tu.1.E (2017) and JLT (2018)]

Maximum Supply Power Constraint

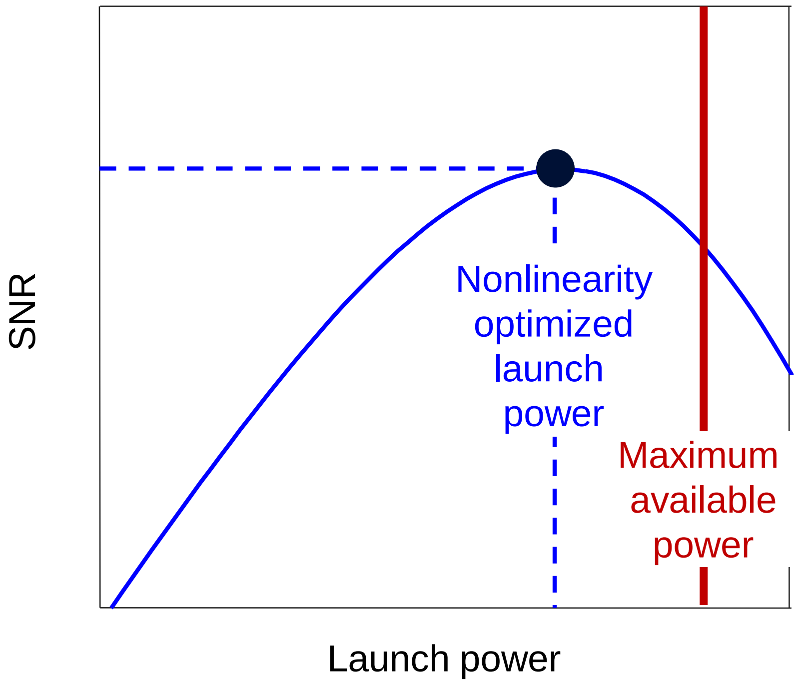

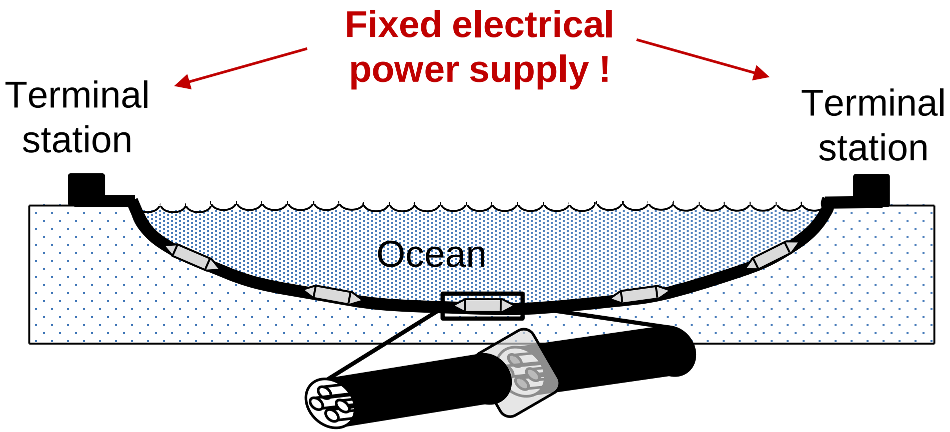

[R. Dar et al., Proc. ECOC, Tu.1.E (2017) and JLT (2018)]

NOKIA

Maximum Supply Power Constraint

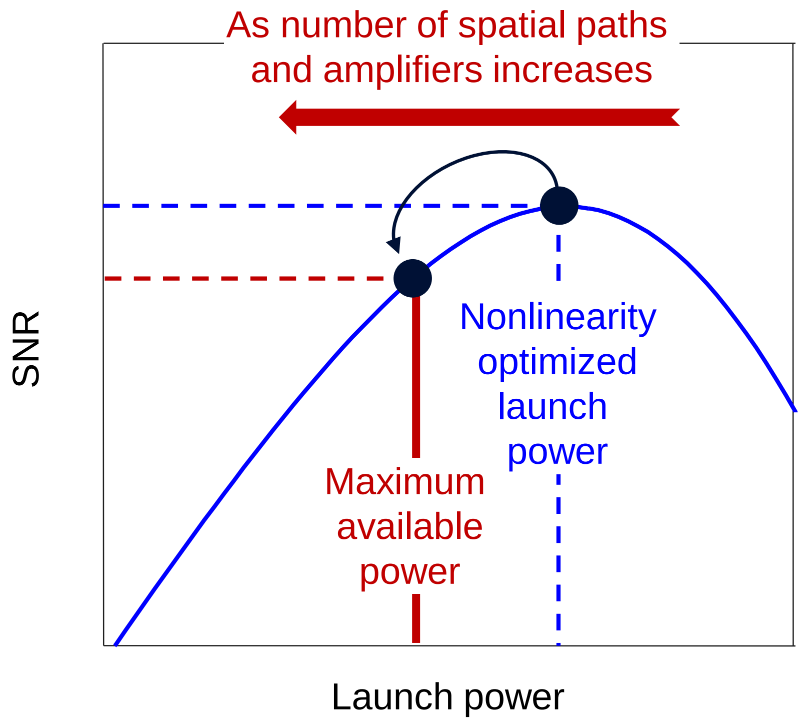

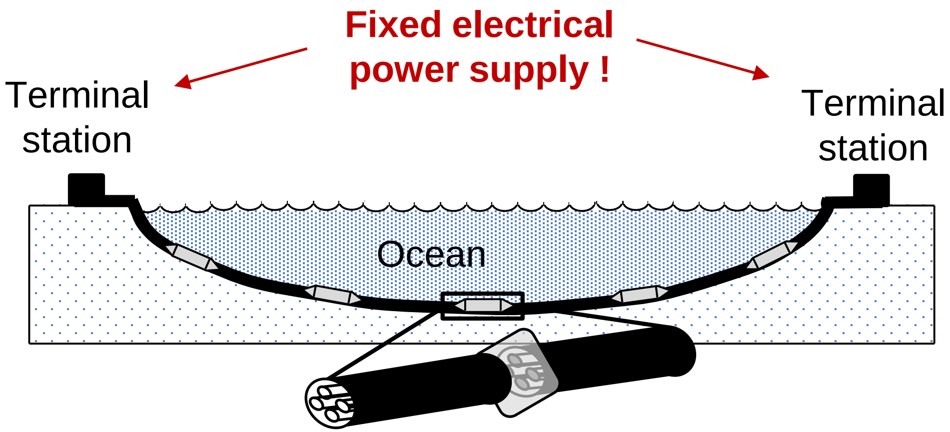

[R. Dar et al., Proc. ECOC, Tu.1.E (2017) and JLT (2018)]

NOKIA

Implications of a Cost Optimized Submarine SDM System

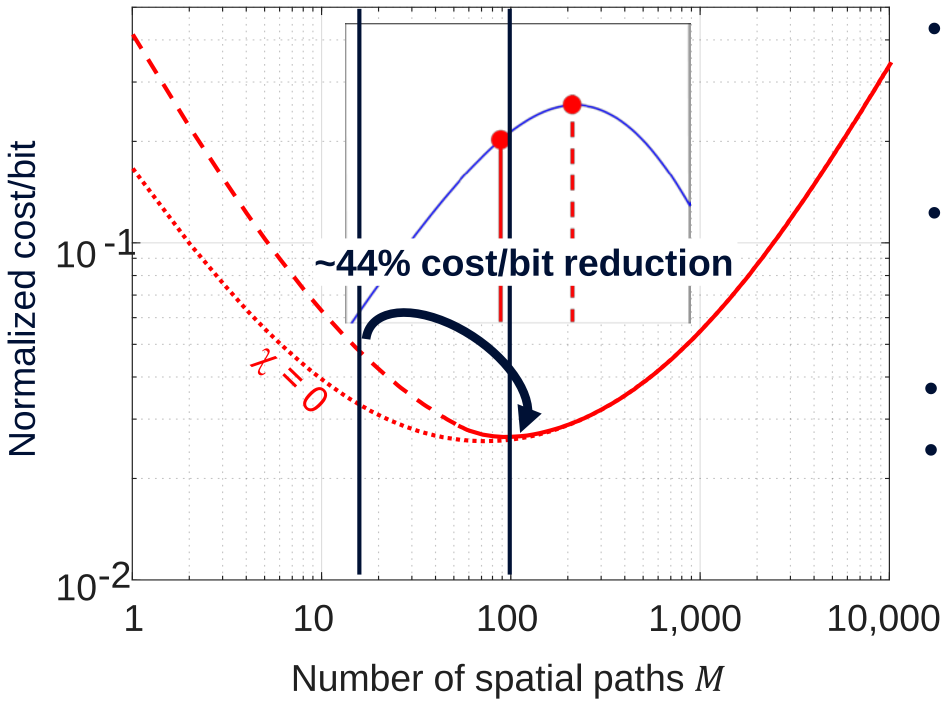

- Nonlinearities become insignificant
  - Low-NL fiber becomes less relevant
  - Digital NL becomes less relevant
- Significant cost/bit savings for ~100 fibers per cable (even without any integration!)

- SDM integration may save another ~35%
- SDM fiber could sell at a premium to avoid higher cabling and deployment costs

[R. Dar et al., Proc. ECOC, Tu.1.E (2017) and JLT (2018)]

NOKIA

Trading Modulation for Multiplexing

A Good Strategy for a Power-Limited Channel

Where should I put my power?

$$C = 2\,M\,B\,log_2(1 + SNR)$$

M and B are ***not*** equivalent, as amplifier gain flattening means killing power !

[A Pilipetskii., OFC Tutorial (2015)]
[O. Sinkin, Phot. Technol. Lett. (2017)]

Implications of a Cost Optimized Submarine SDM System

	Half C band	C band	C+L band
Bandwidth	18 nm	35 nm	70 nm
OA efficiency	6.5%	2.5%	1.3%
Noise figure	5.4 dB	5.0 dB	5.7 dB
Normalized cost	0.7	2	3.6
System cost/bit	0.0226	0.0268	0.0305
Optimum M	350	90	34
Cable Capacity	4.96 Pb/s	2.37 Pb/s	1.52 Pb/s

$$C = MB\log_2\left(1 + \frac{\eta_{OE}P_e}{N^2 MBhf\left(Fe^{\frac{\alpha L}{N}} - 1\right)}\right)$$

- Amplifier bandwidth x Spatial paths
- Amplifier efficiency / Noise figure
- More supply power doesn't help much

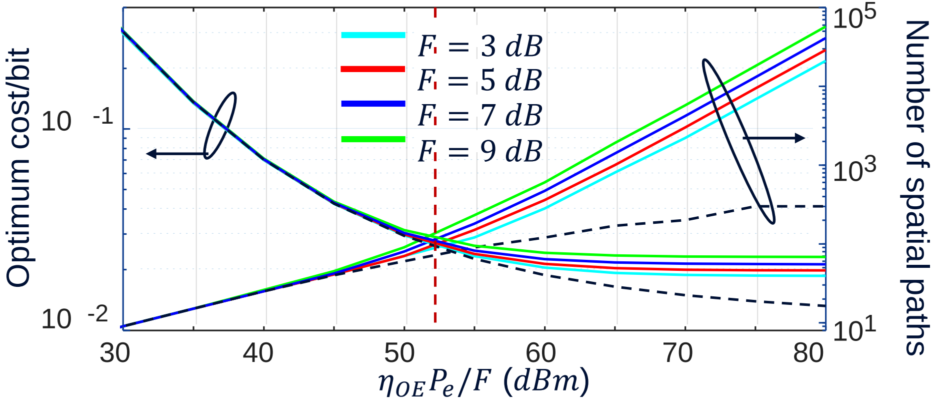

[R. Dar et al., Proc. ECOC, Tu.1.E (2017) and JLT (2018)]

Bandwidth scaling is only required if parallel fiber is not available or too expensive to deploy (which is unfortunately frequently the case)

$$C = 2\, M\, B\, log_2(1 + SNR)$$

Polarization | Bandwidth

Spatial paths

Pre-log (multiplexing) factors   Logarithmic (modulation) capacity

NOKIA

Ultra-Wideband Amplified Systems

$C \sim B \times \log_2(SNR)$

100-nm SOA

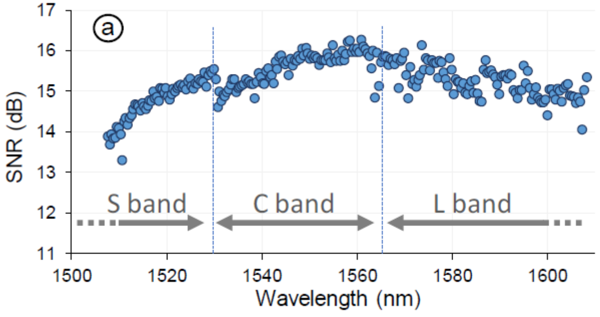

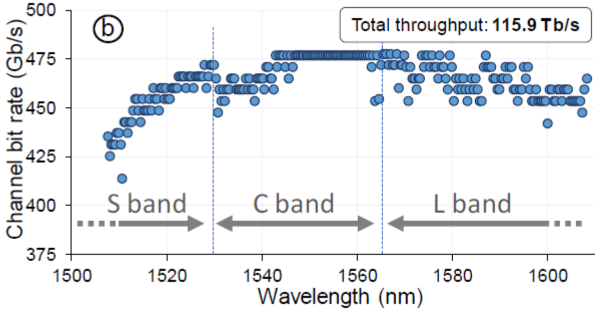

[J. Renaudier et al., Proc. ECOC, Th.PDP.A.3 (2017)]

160-nm Raman amplifier in Tellurite fiber

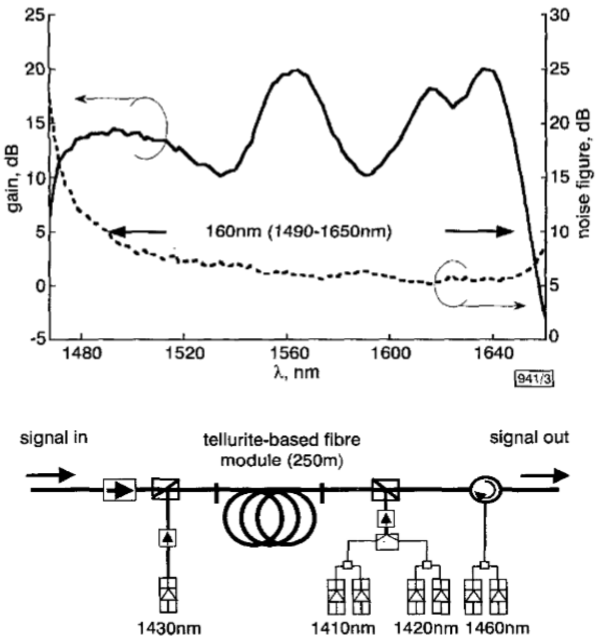

[A. Mori et al., El. Lett. 1442 (2001)]

Be Careful With Bandwidth Scaling Phantasies

$C \sim B \times \log_2(SNR)$

What Counts in Engineering is the *Relative* Bandwidth

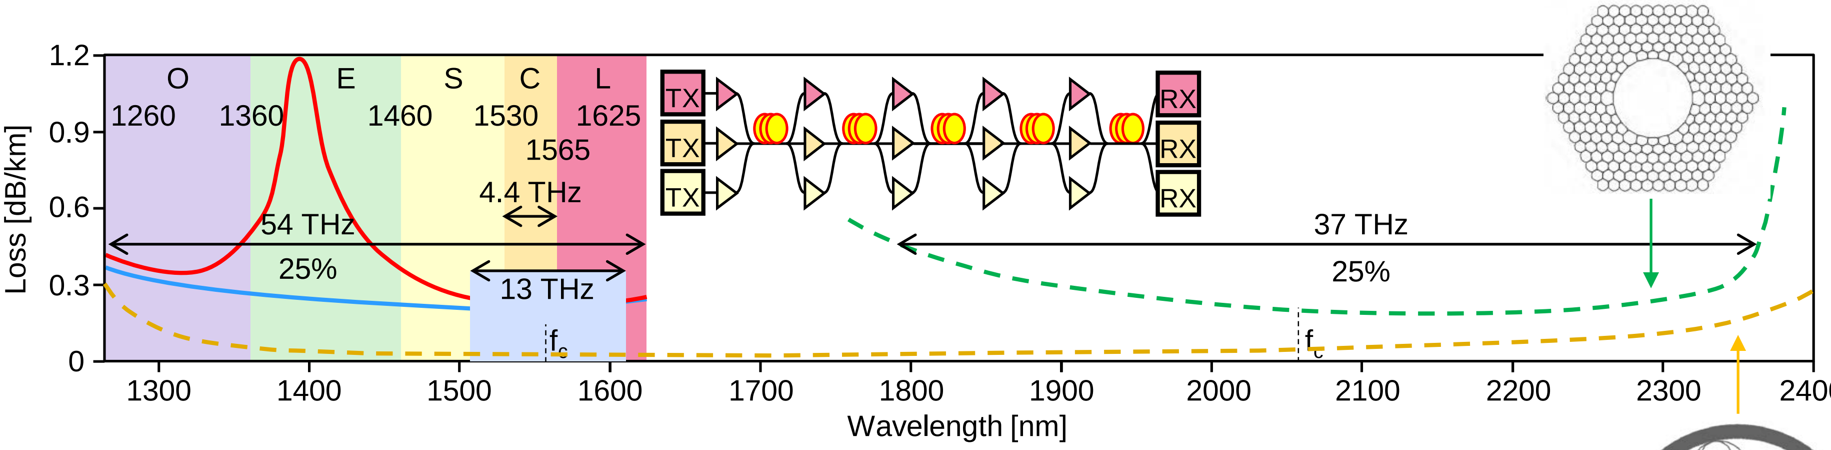

Hollow-core fiber (HCF)

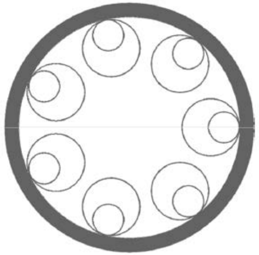

Relative BW = $\frac{\text{System BW}}{\text{Center frequency}}$

C-band: 2.3%

100-nm SOA: 6.6% [Renaudier et al., ECOC 2017]

NAN-HCF: 67% (= octave-spanning)

Nested antiresonant nodeless HCF

[D. J. Richardson, tutorial Tu3H.1, OFC 2017]

NOKIA

Absolute Bandwidth, Relative Bandwidth, and Carrier Frequency

Is Going to the Extreme UV or the Soft X-Ray Range a Crazy Idea ?

(In analogy to the transition from electrical cables to optical fiber in the 1970s)

$$C = 2\,M\,B\,log_2(1 + SNR)$$

$$C = 2\,M\,\underbrace{f_c\,B_{rel}}_{B}\,log_2(1 + SNR)$$

Polarization

Spatial paths

Carrier frequency

Relative bandwidth

Pre-log (multiplexing) factors

Logarithmic (modulation) capacity

NOKIA

Absolute Bandwidth, Relative Bandwidth, and Carrier Frequency

Is Going to the Extreme UV or the Soft X-Ray Range a Crazy Idea ?

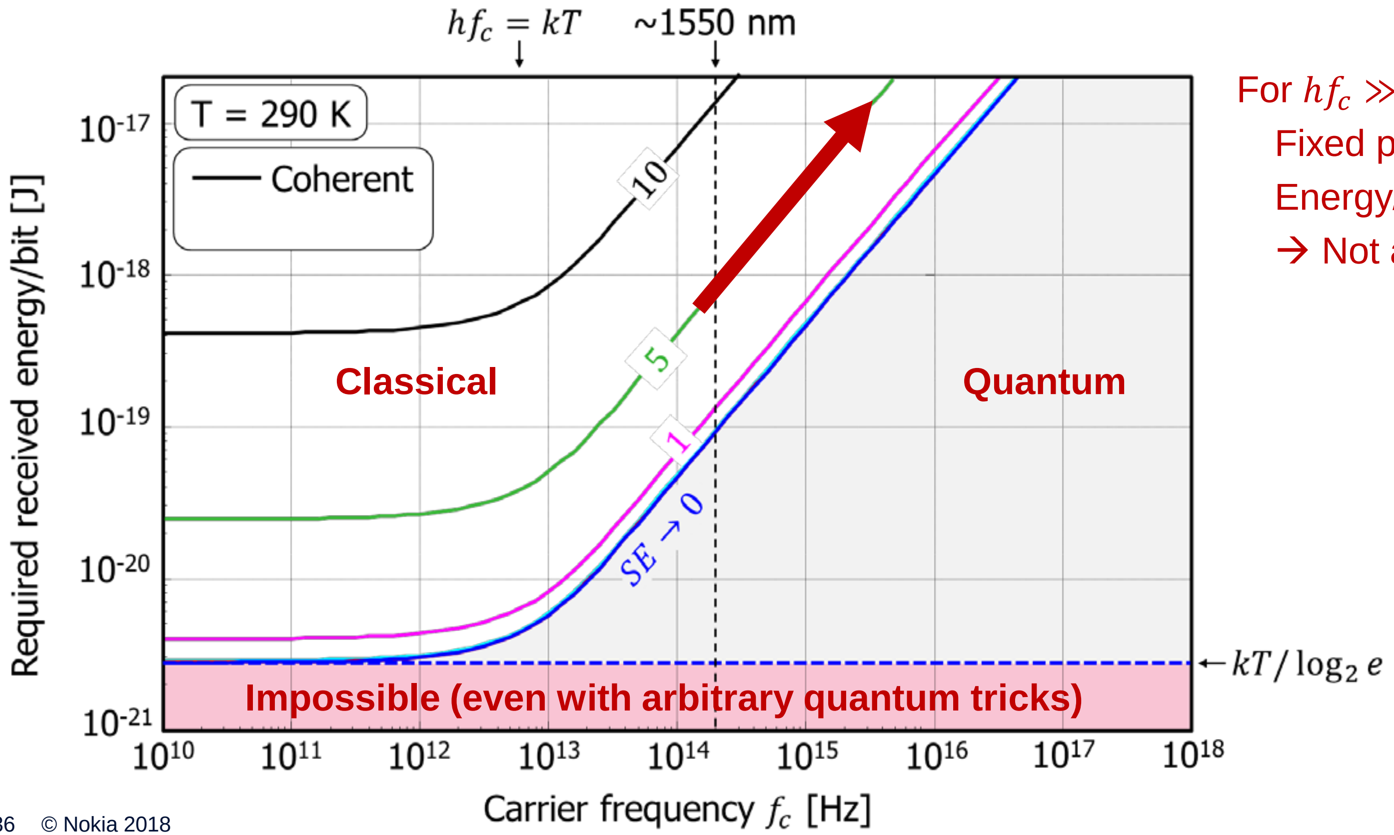

For $hf_c \gg kT$:

- Fixed photons/bit at fixed SE
- Energy/bit ~ hf_c
- → Not a viable scaling

Absolute Bandwidth, Relative Bandwidth, and Carrier Frequency

Is Going to the Extreme UV or the Soft X-Ray Range a Crazy Idea ?

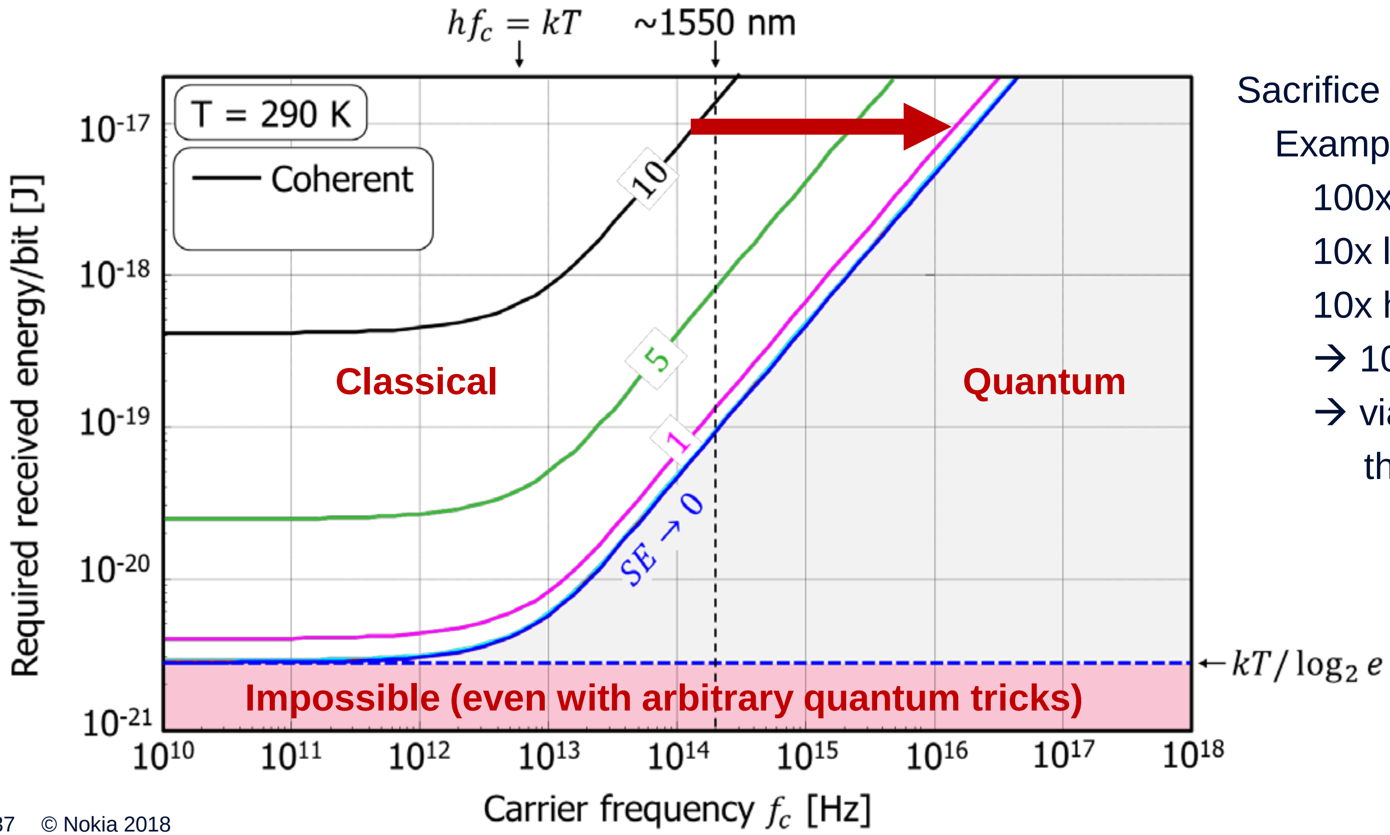

Sacrifice SE while increasing M

- Example:
  - 100x higher f_c (100x more C)
  - 10x lower SE
  - 10x higher M
  - → 10x net gain at constant M
  - → viable scaling, even if less than hoped for

Absolute Bandwidth, Relative Bandwidth, and Carrier Frequency

Is Going to the Extreme UV or the Soft X-Ray Range a Crazy Idea ?

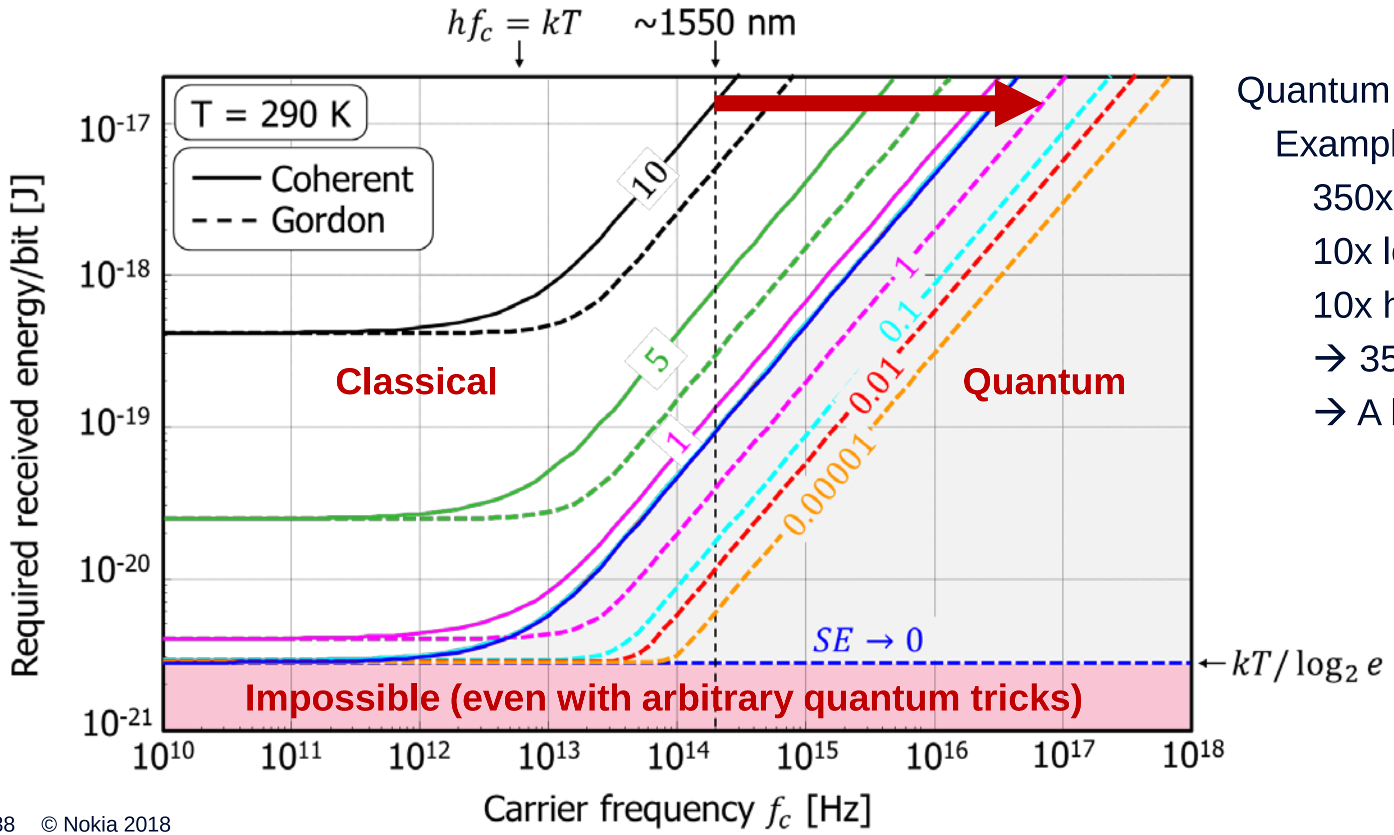

Quantum techniques (Gordon)

- Example:
  - 350x higher f_c (350x more C)
  - 10x lower SE
  - 10x higher M
  - → 35x net gain at constant M
  - → A bit better than classical

Fairly Limited Capacity Gains When Going to Higher Carrier Frequencies

Requires as of Yet Unknown Quantum and UV/X-Ray Technologies

→ Probably Indeed a Crazy Idea

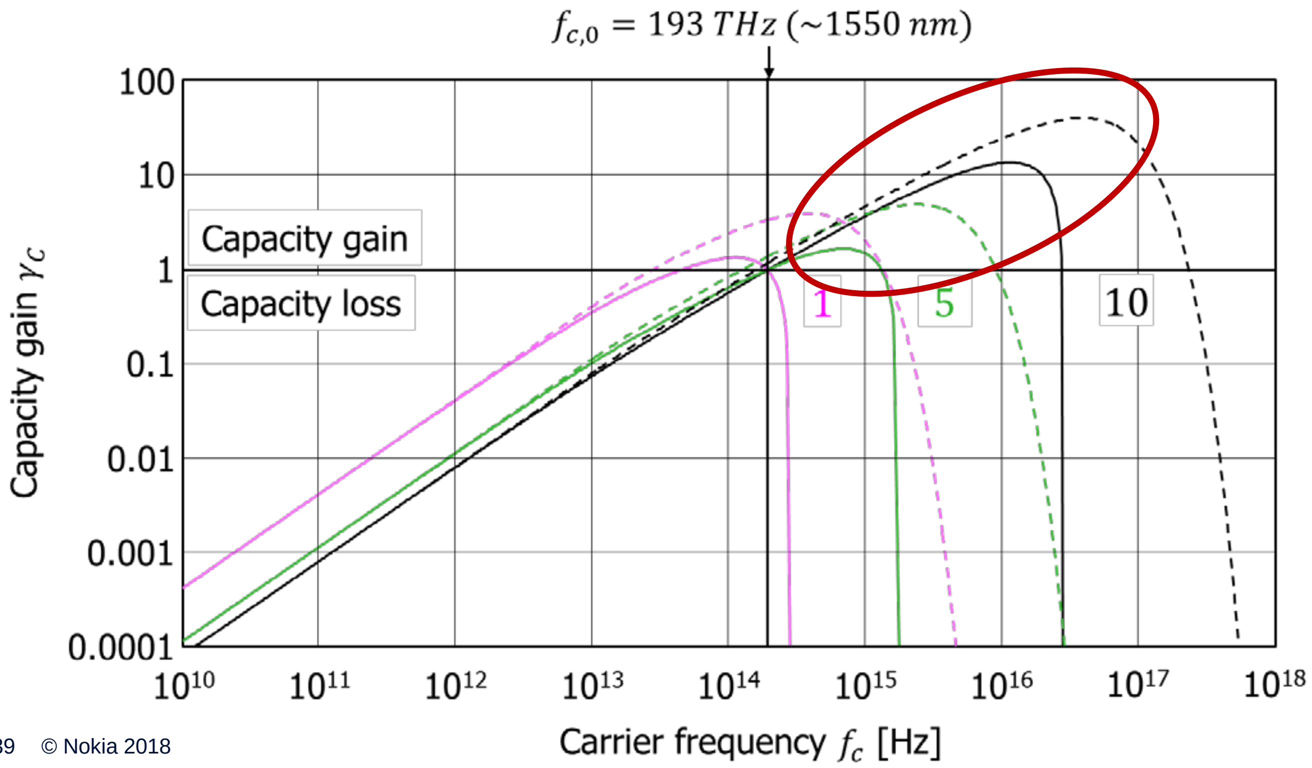

And the winner is…

$$C = 2\, M\, f_c\, B_{rel}\, log_2(1 + SNR)$$

Polarization

Spatial paths

Carrier frequency

Relative bandwidth

Pre-log (multiplexing) factors

Logarithmic (modulation) capacity

Full Parallelism Leads to WDM x SDM Systems

What Might a 10T Interface in a 1P System Look Like in ca. 2024?

WDM x SDM

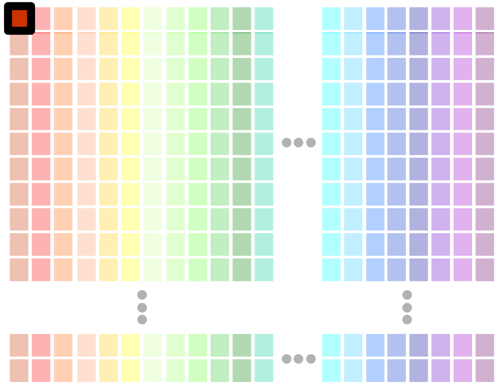

- Matrix of “unit cells” in WDM x SDM space → *Replicate simple unit cells*
- *Bandwidth* of unit cell driven by high-speed opto-electronics
- *Bit rate* of unit cell driven by symbol rate and modulation format

Inside a unit cell: High-speed modulation records

107 GBd 16-QAM **(856 Gb/s)** | 72 GBd 64-QAM **(864 Gb/s)** | 90 GBd 64-QAM **(1.08 Tb/s)**

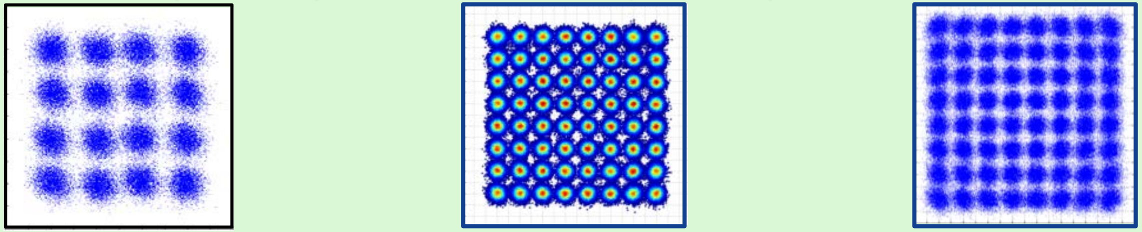

[G. Raybon et al., ECOC 2013] | [S. Randel et al., OFC 2014] | [G. Raybon et al., IPC 2015]

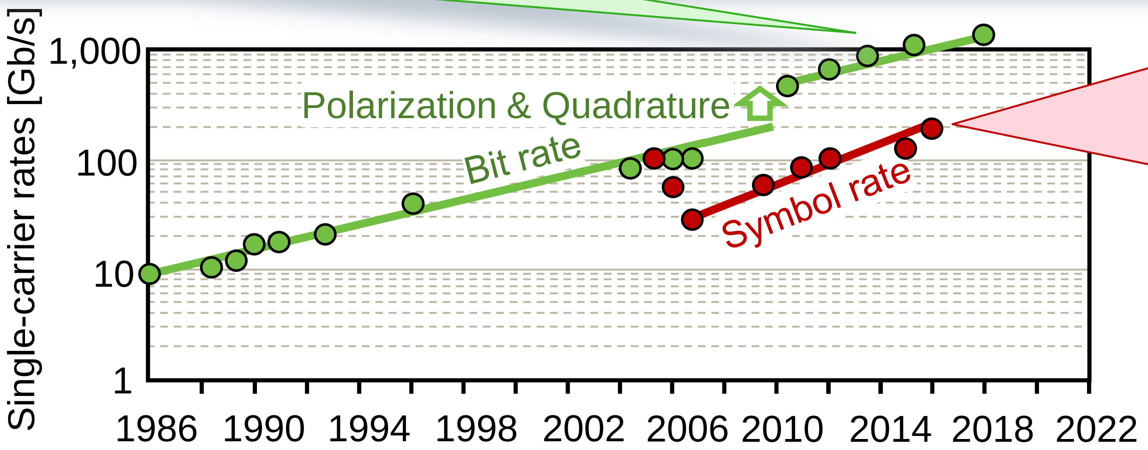

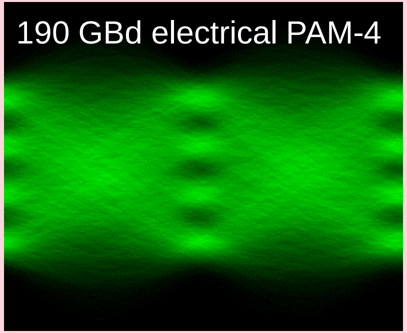

[X. Chen et al., OFC 2016]

Pushing Interface Rates to New Limits

Commercial Perspective of High-Speed Opto-Electronics & DSP

- Commercial CMOS ASICs with converters for symbol rates of 100^{+} Gbaud
- Commercial 1T interfaces with meaningful transmission reach
- Commercial multi-core ASIC DSP processing power of several (10?) Tb/s

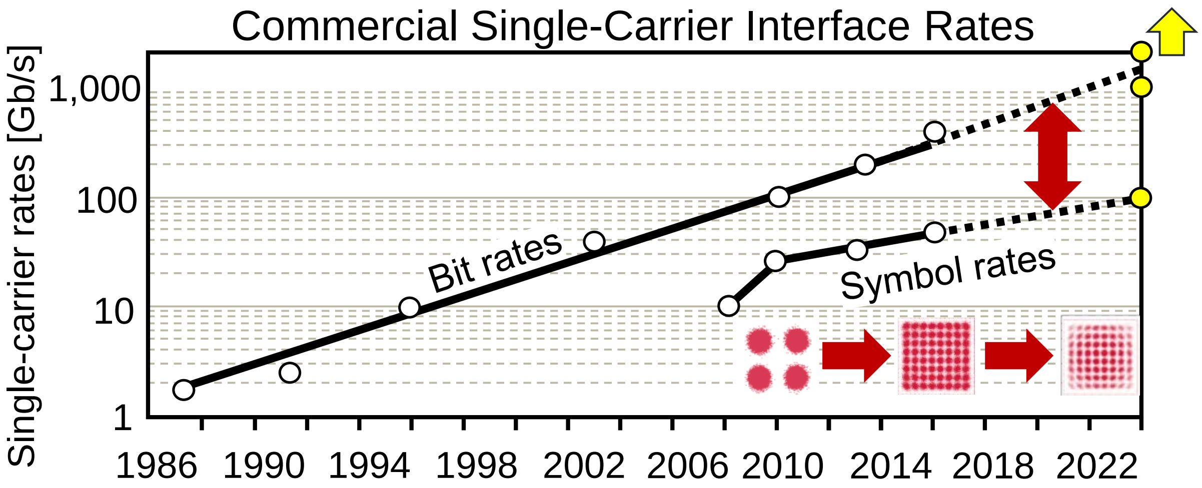

NOKIA

Spectral vs. Spatial Superchannels

Spectral Superchannel

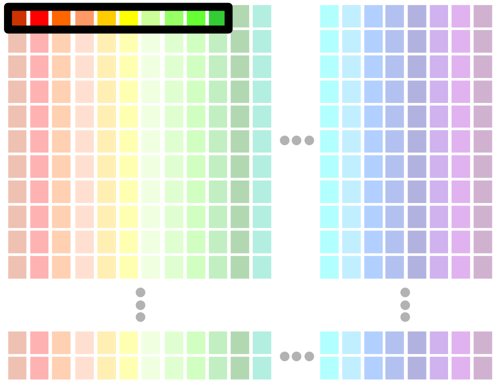

Traditional technology
Pay as you grow
Filtering and switching advantages
Compensate cross-channel NL

Spatial Superchannel

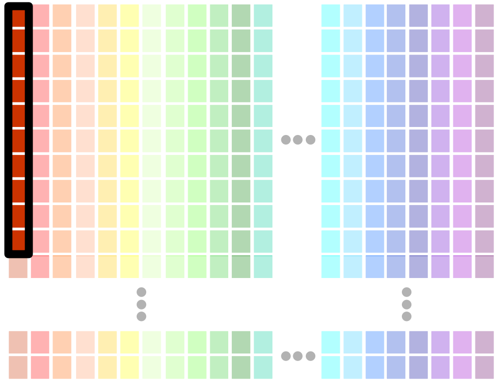

Parallel paths from day one
Better array integration
Compensate array crosstalk

Metro / Long-haul Networking ↔ ***Submarine & DCI Point-to-Point***

Spectral superchannels — Spatial superchannels

NOKIA

Using SDM, We Can Comfortably Scale Networks for the Next 20 Years

Table 2. Possible system evolutions over the next 10 and 20 years.

	2017	2027	2037
Symbol rate [GBaud]	50	120	300
Bit rate [Gb/s]	200 – 400	600 – 1,600	2,000 – 6,000
System bandwidth [THz]	5	5 – 12	5 – 20
Capacity per spatial path [Tb/s]	20 – 40	25 – 160	32 – 400
Unit cells per spatial path	100	40 – 100	16 – 66
Target system capacity [Pb/s]	0.02 – 0.04	1 – 2	50 – 100
Required number of spatial paths	1	6 – 80	125 – 3125

Fiber-Optic Transmission and Networking: The Previous 20 and the Next 20 Years

PETER J. WINZER*, DAVID T. NEILSON, AND ANDREW R. CHRAPLYVY

Nokia Bell Labs, 791 Holmdel Road, Holmdel, NJ 07733, USA
peter.winzer@nokia-bell-labs.com

NOKIA

Like in Any Parallel Solution: Volume is Good, But Integration is Key

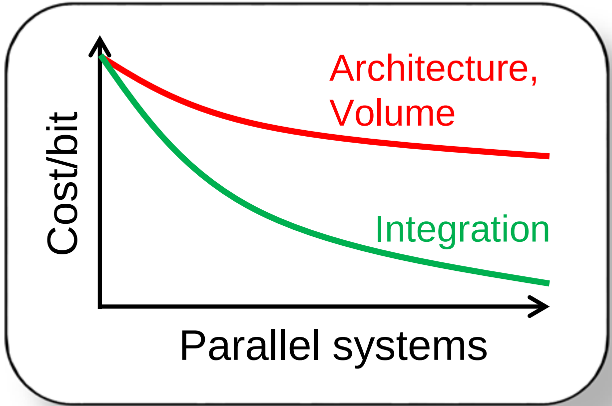

- Multi-channel modulators
- Multi-channel drivers
- Multi-channel receivers
- Multi-channel ASICs
- Multi-channel optical amplifiers
- Parallel optical switch elements
- Multi-path fibers, connectors, splices
- Shared power supplies, comb sources, etc.

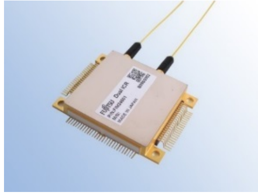

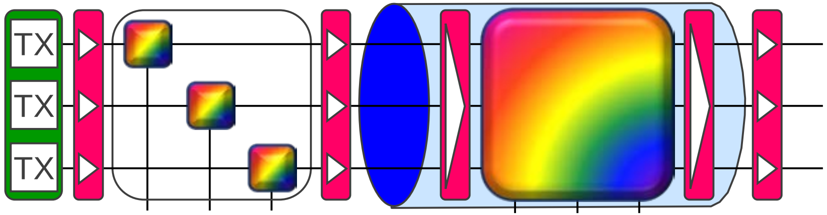

Integration of Parallel Fiber Channels

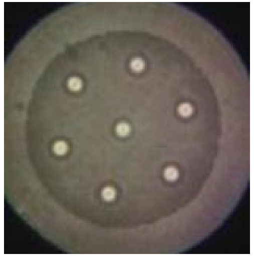
[Zhu, ECOC 2011]

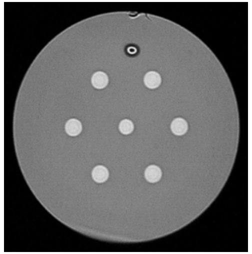
[Hayashi, ECOC 2011]

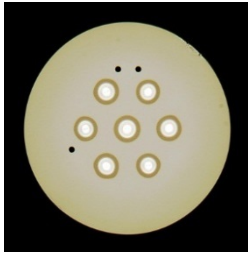
[Imamura, ECOC 2011]

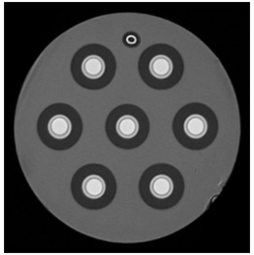
[Hayashi, OFC 2011]

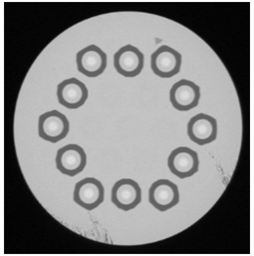
[Takara, ECOC 2012]

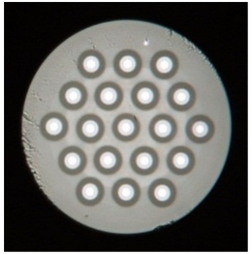
[Sakaguchi, OFC 2012]

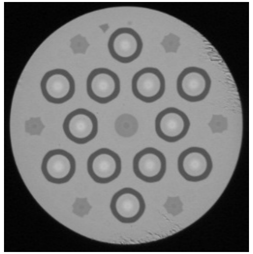
[Kobayashi, ECOC 2013]

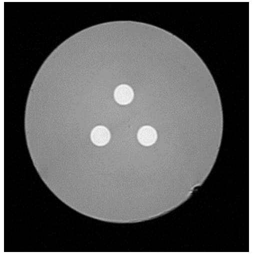
[Ryf, ECOC 2011]

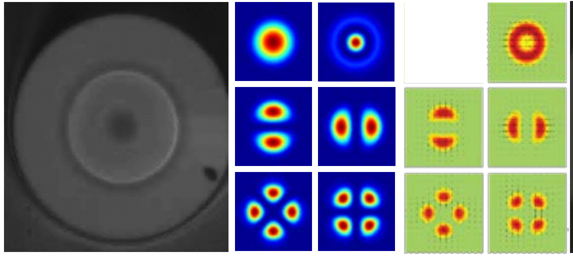
[Ryf, FiO 2012]

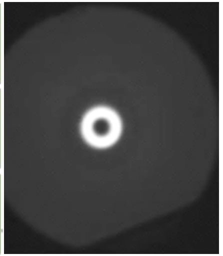
[Doerr, ECOC 2011]

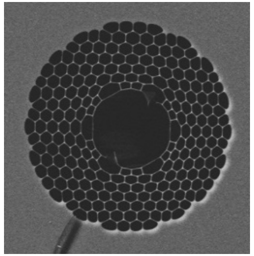
[Petrovich, ECOC 2012]

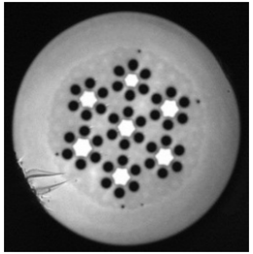
[Cia, IPS SumTop 2012]

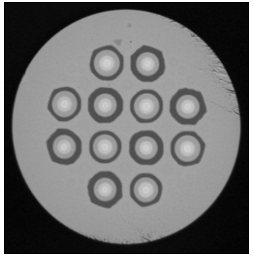
[Mizuno, OFC 2014]

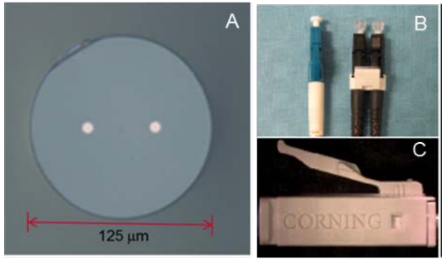

[Y. Geng et al., SPIE 2015]

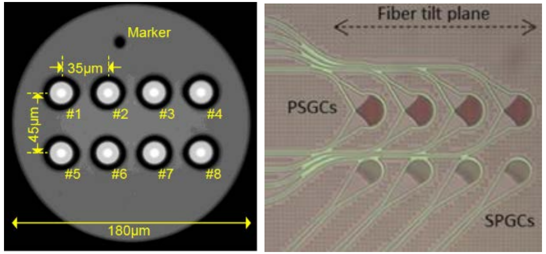

[T. Hayashi et al., ECOC 2017]

Mode Coupling Increases Nonlinear Transmission Performance
Long-Haul Experimental Confirmation

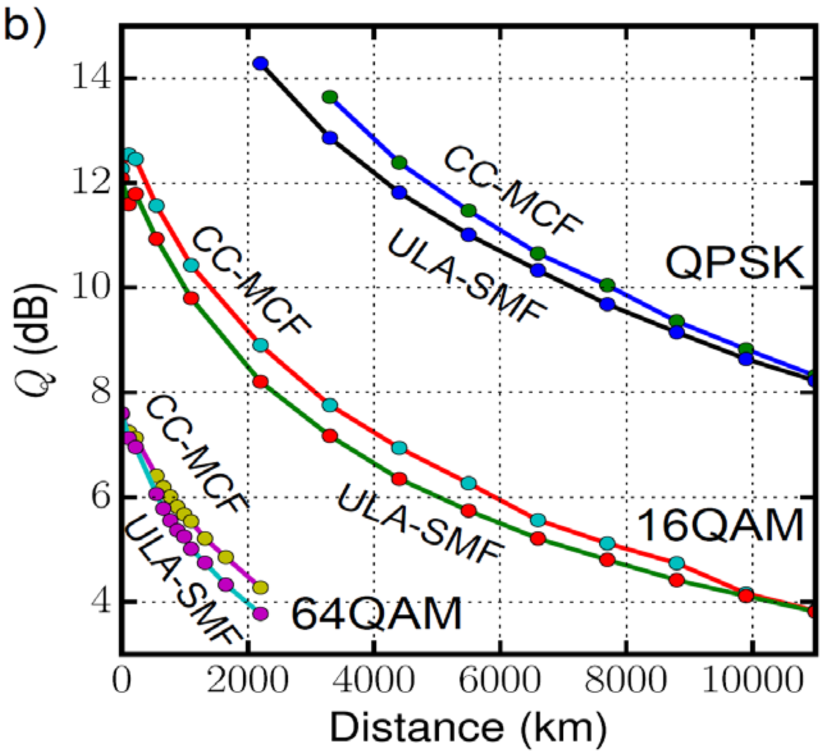

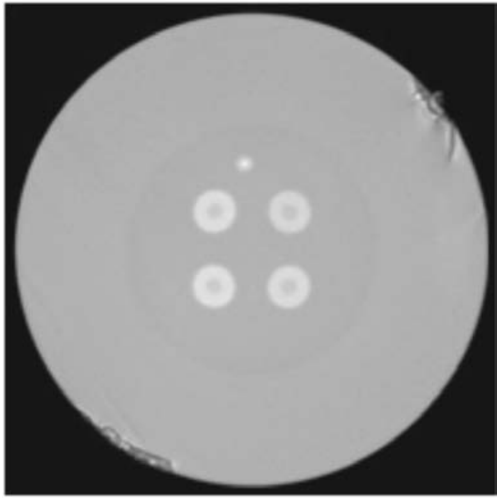

[R. Ryf et al., ECOC-PD (2016)]

NOKIA

MIMO for SDM is Not a Big Problem

ASICs are Limited by CD Filter and FEC

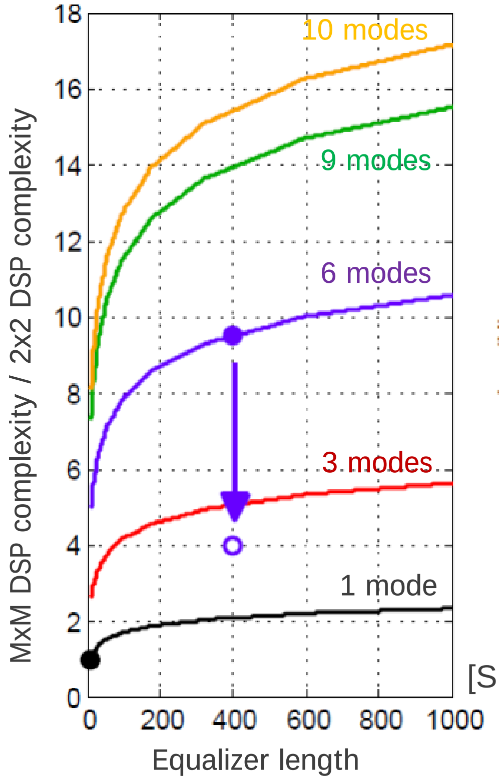

- 4x MIMO DSP complexity for 6 modes (2x2 MIMO → 12x12 MIMO)
- MIMO is only ~10% of overall DSP today:
  ➔ Only ~1.3x higher ASIC complexity

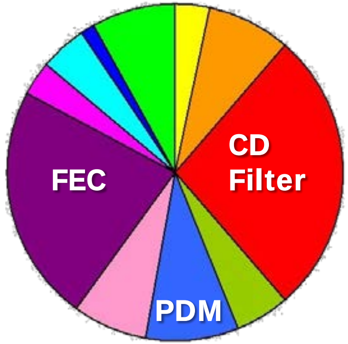

- The real problem: Interfacing of many coherent frontends

[S. Randel et al., ECOC 2013, Th.2.C.4]

Three Key Aspects of Transponder Integration

1. The Opto-Electronic Array Integration Challenge

4 x 25G (PSM4)
[Y. De Koninck, ECOC 2017]

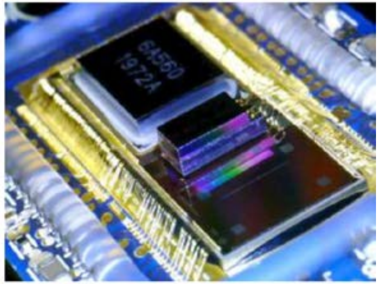

16 x 25G
[T. Aoki, ECOC 2017]

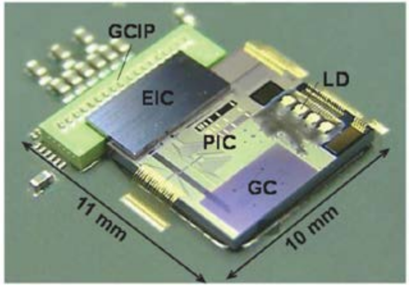

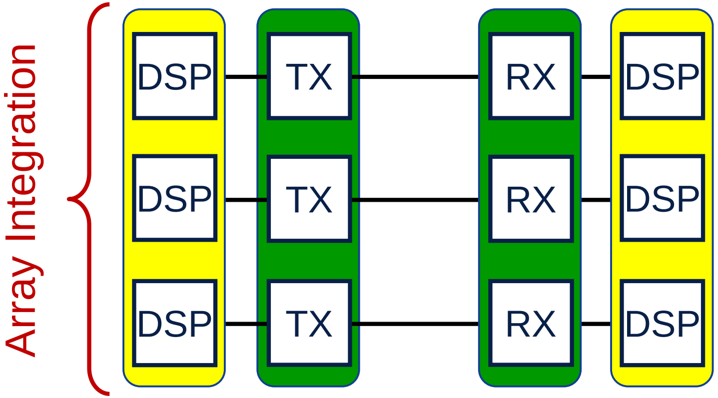

- **10T** = 10 x 1T = 100 x 100G
- Reduced speed ➔ Higher parallelization
- Non-negotiable: Cost, energy, footprint
- Similar to short-reach array approaches

NOKIA

Three Key Aspects of Transponder Integration

2. The Optics-Electronics Array Integration Challenge

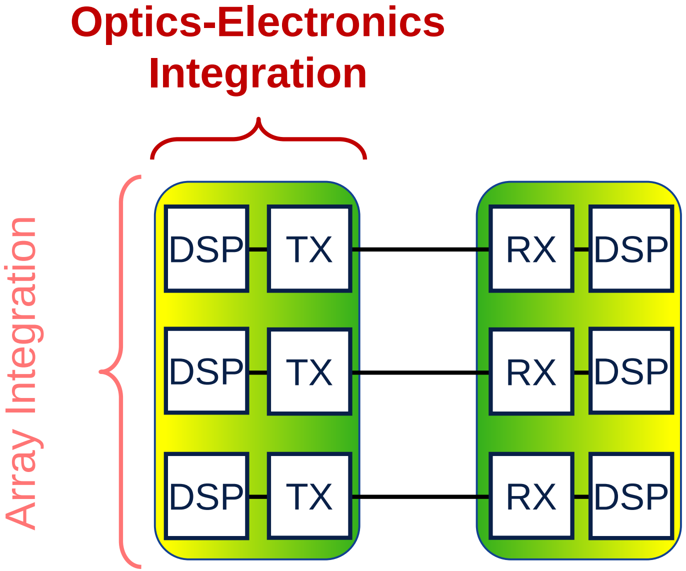

Courtesy: M. Smit, K. Williams

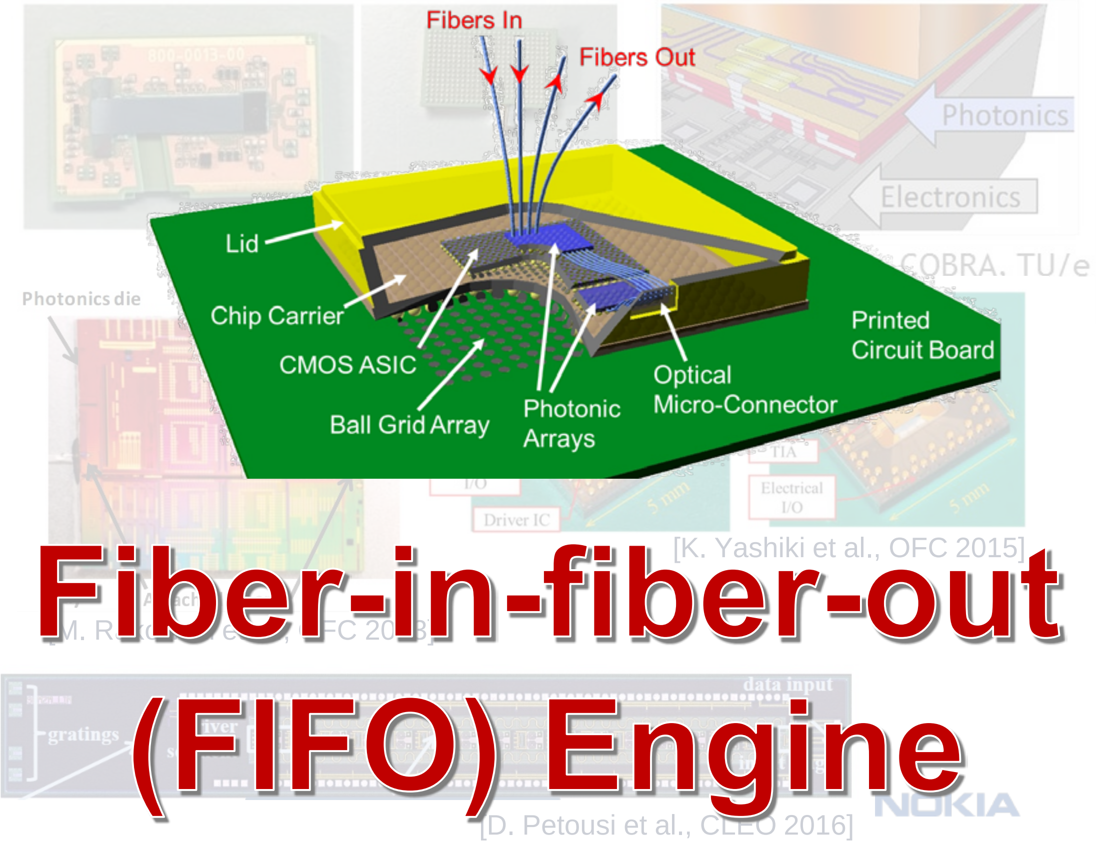

Fiber-in-fiber-out (FIFO) Engine

Three Key Aspects of Transponder Integration

3. The Holistic Integration of Opto-Electronics With DSP

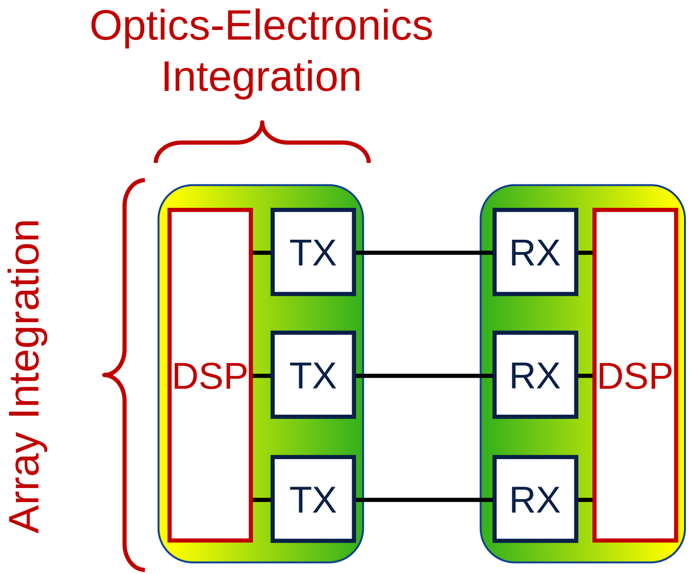

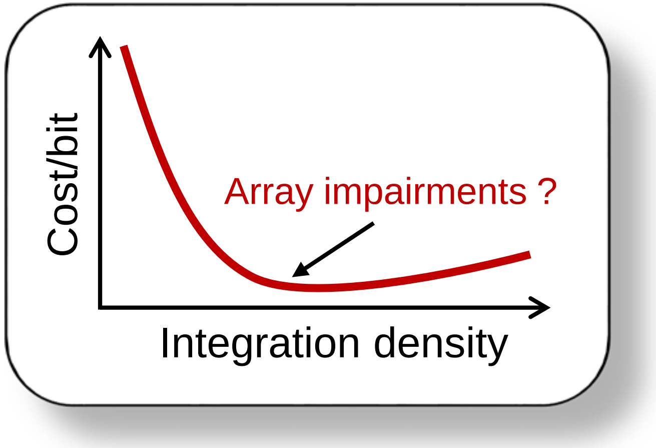

Remove array-impairments through DSP (MIMO, etc.)

NOKIA

Three Key Aspects of Transponder Integration

3. The Holistic Integration of Opto-Electronics With DSP

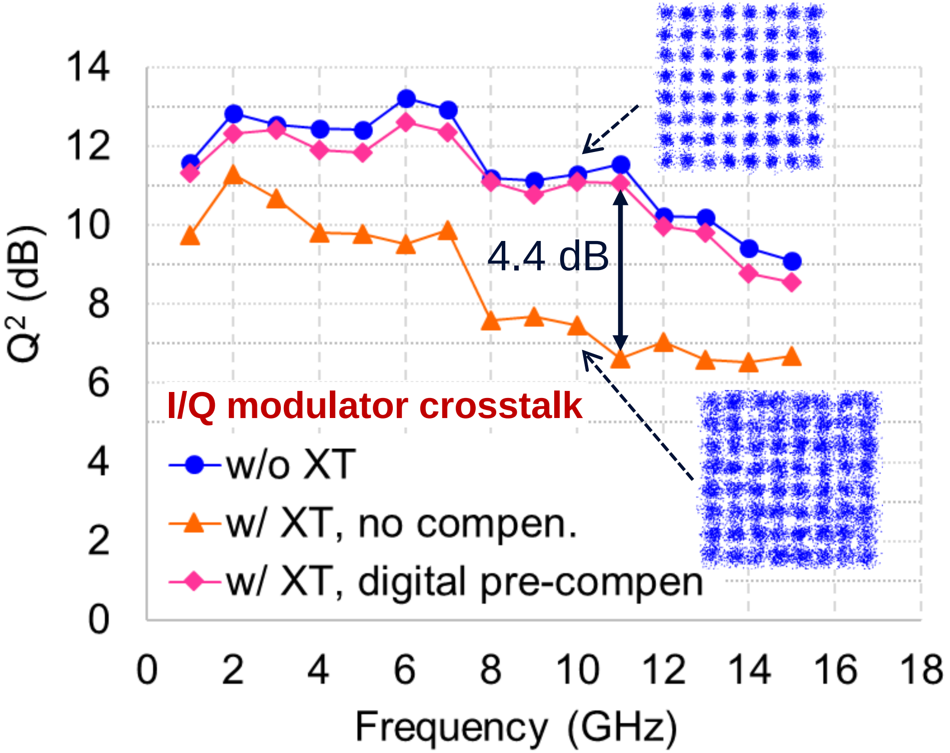

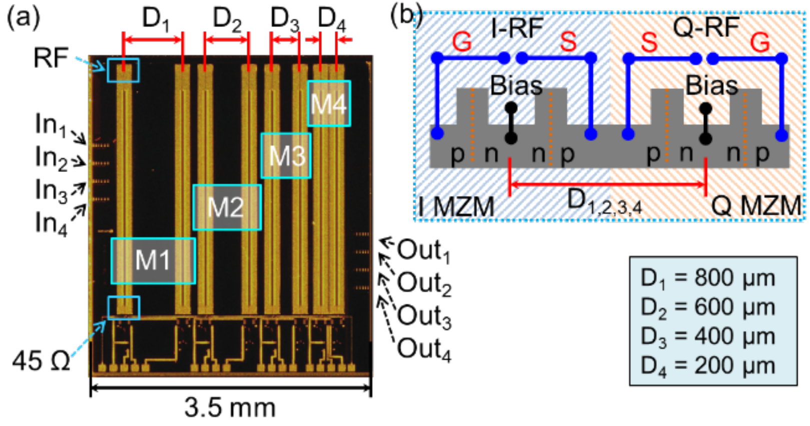

[X. Chen et al., ECOC 2016]

Fiber Tapping

What's the Role of MIMO-SDM ?

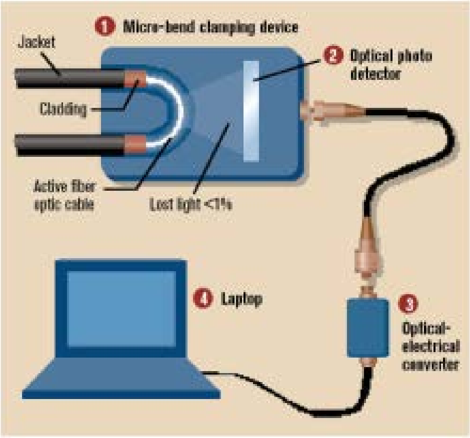

Sandra Kay Miller,
Information Security Magazine,
November 2006

- In SDM: Mode dependent loss (MDL) from eavesdropping
  - Detect an eavesdropper
  - Achieve provable security against physical layer attacks
- How secure can an SDM waveguide be?
- How can an SDM waveguide be wire-tapped?

- Information-theoretic security metrics (*fundamental* security)
  - "Secrecy Capacity": $C_{Alice\text{-}Bob} – C_{Alice\text{-}Eve}$
  - Force Eve to induce enough MDL or not get enough signal

MIMO-SDM Can Support Very Large “Fundamentally Secure” Bit Rates

A Single Spatial Path is ~10 to 100 Tb/s ! ➔ 7% is 70 Gb/s to 7 Tb/s

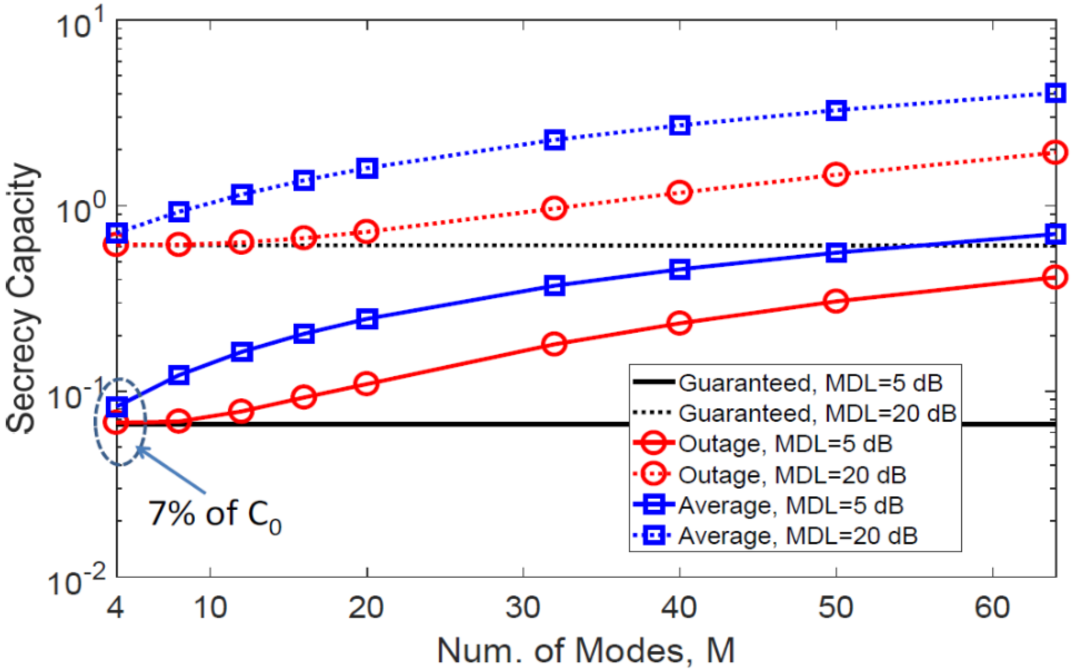

[K. Guan et al., Proc. ECOC 2012]
[K. Guan et al., Asilomar 2012]
[K. Guan et al., IEEE Trans. Inf. Sec. For. (2015)]
[K. Guan et al., Opt. Commun. (2018)]

NOKIA

Conclusions

1. Network traffic growth remains strong
   Sold WDM transponder capacity scales at ~45%
   Scaling is widely supported by compute, storage, access scaling
   Optical transport is significantly falling behind
   → Worrisome scaling disparities ("Capacity Crunch")
2. Parallelism is mandatory
   log(SNR) scaling vs. linear (pre-log) multiplexing gains
3. By 2025 we will need 10T interfaces in 1P systems
   Massively integrated (coherent) systems: B x M x log(SNR)
   Scaling using WDM x SDM
4. Integration, integration, integration
   Holistic view of client and line interfaces integrated onto digital CMOS chips
   "Fiber-in-fiber-out" (FIFO) all-in-one transport processor solutions
5. Interesting "fundamental security" aspects

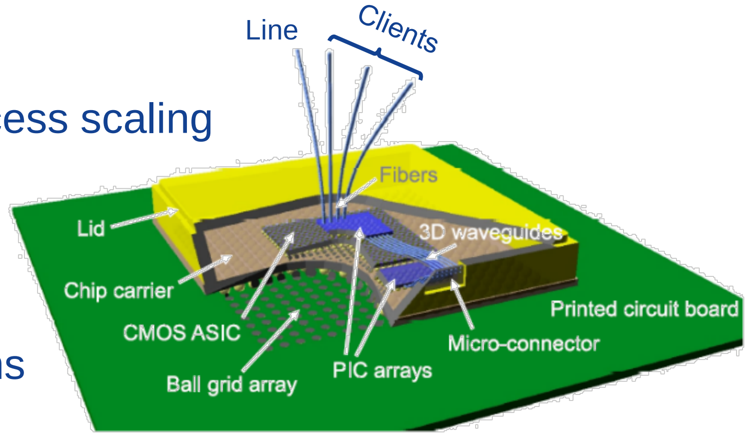

NOKIA


NOKIA