How to Build a Petabit-per-Second Optical Router

Rod Tucker

Centre for Ultra-Broadband Information Networks (CUBIN)

Department for Electrical and Electronic Engineering

University of Melbourne, Australia

Optical Network Evolution

Largest commercial router: ~50 Tbit/s
http://www.cisco.com/en/US/products/ps5842/index.html

Power Consumption: ~1 MW
Footprint: ~100 m^2

Can we do better with optics?

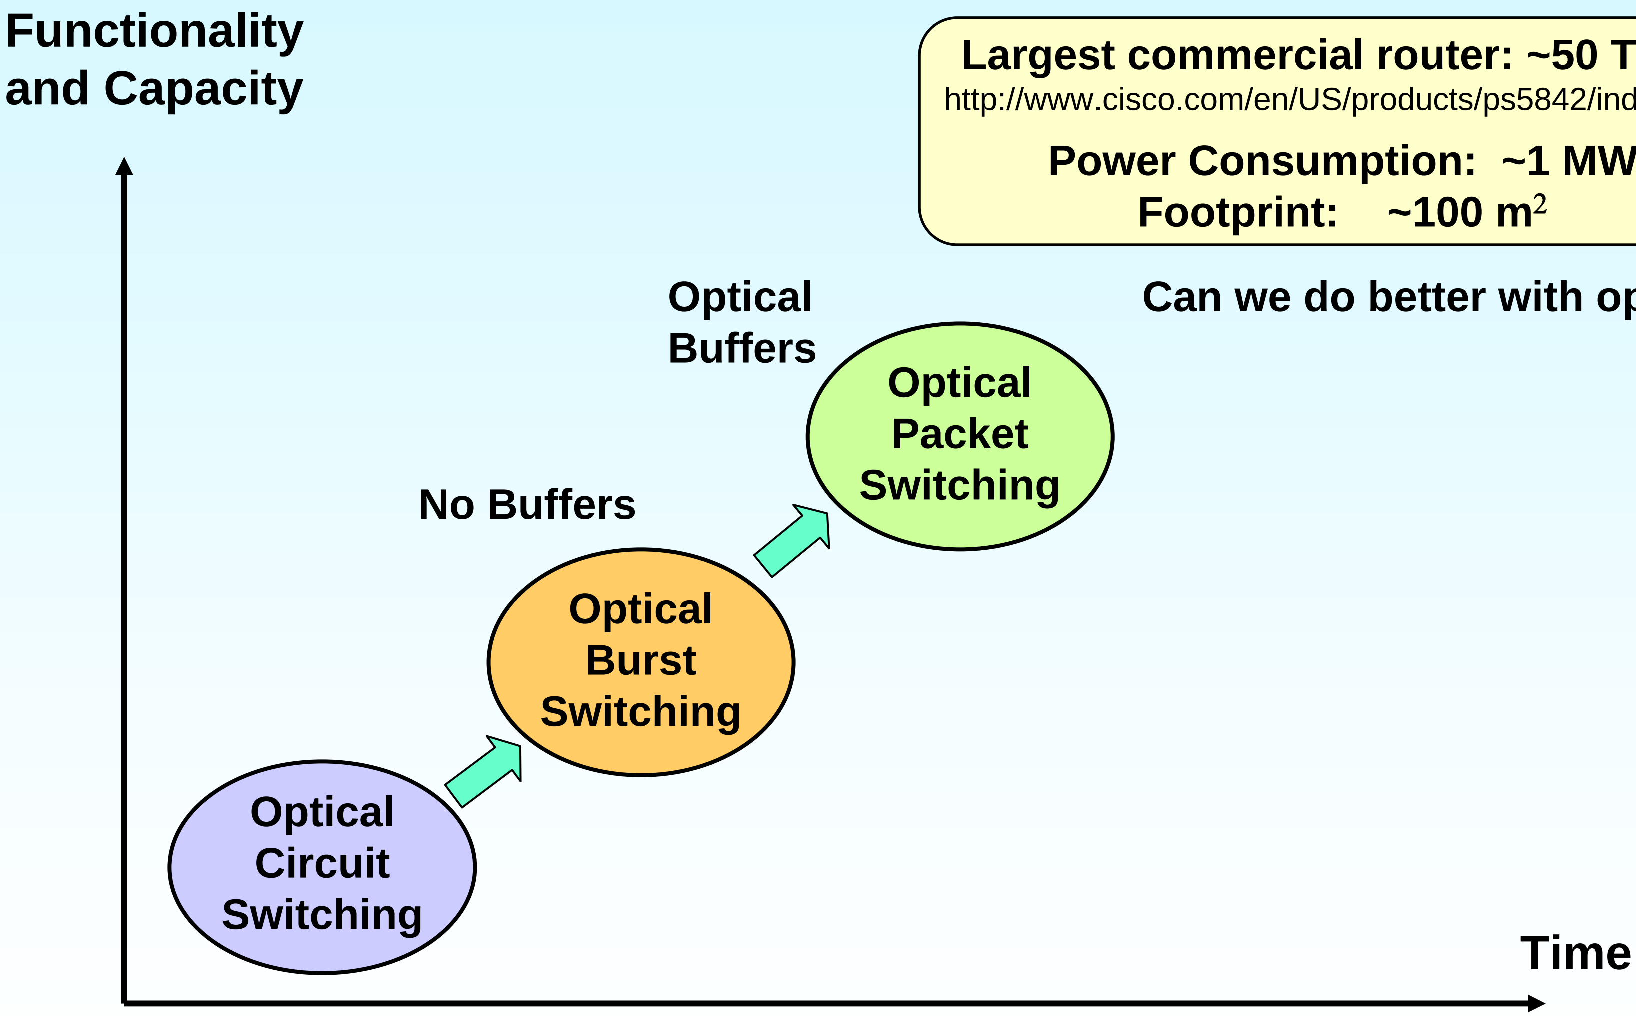

Why Petabit-per-Second Routers

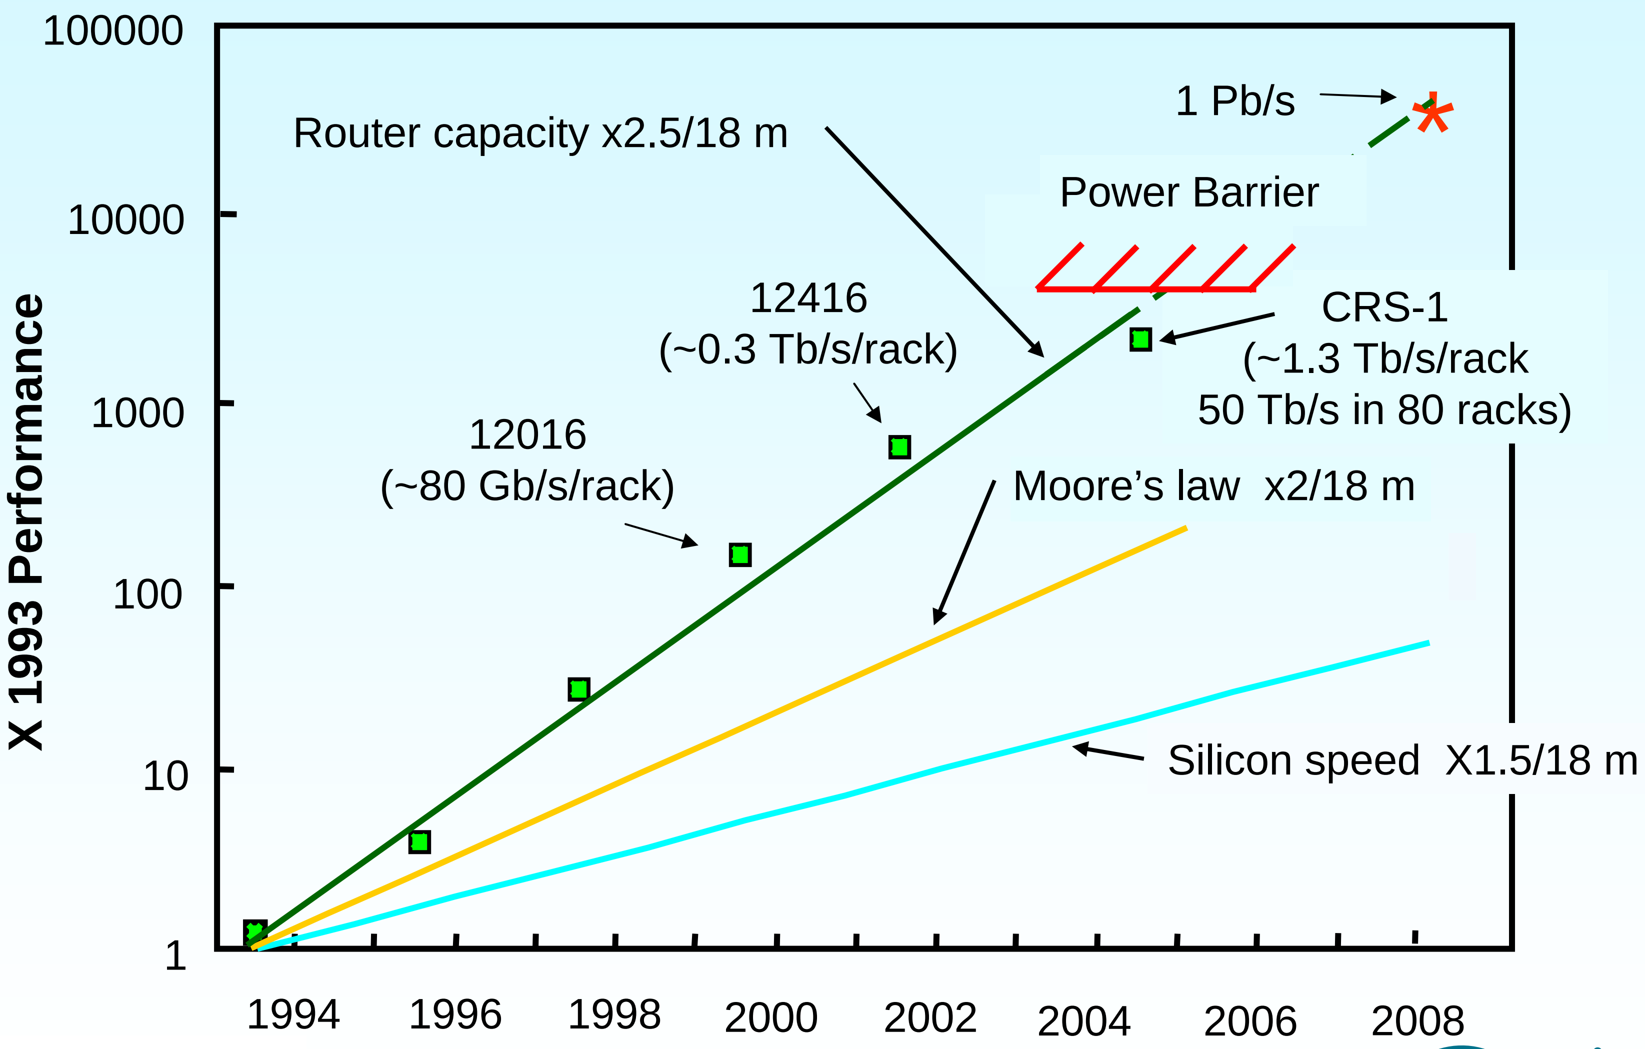

Source: G. Epps, CISCO

Power Consumption in Electronic Routers

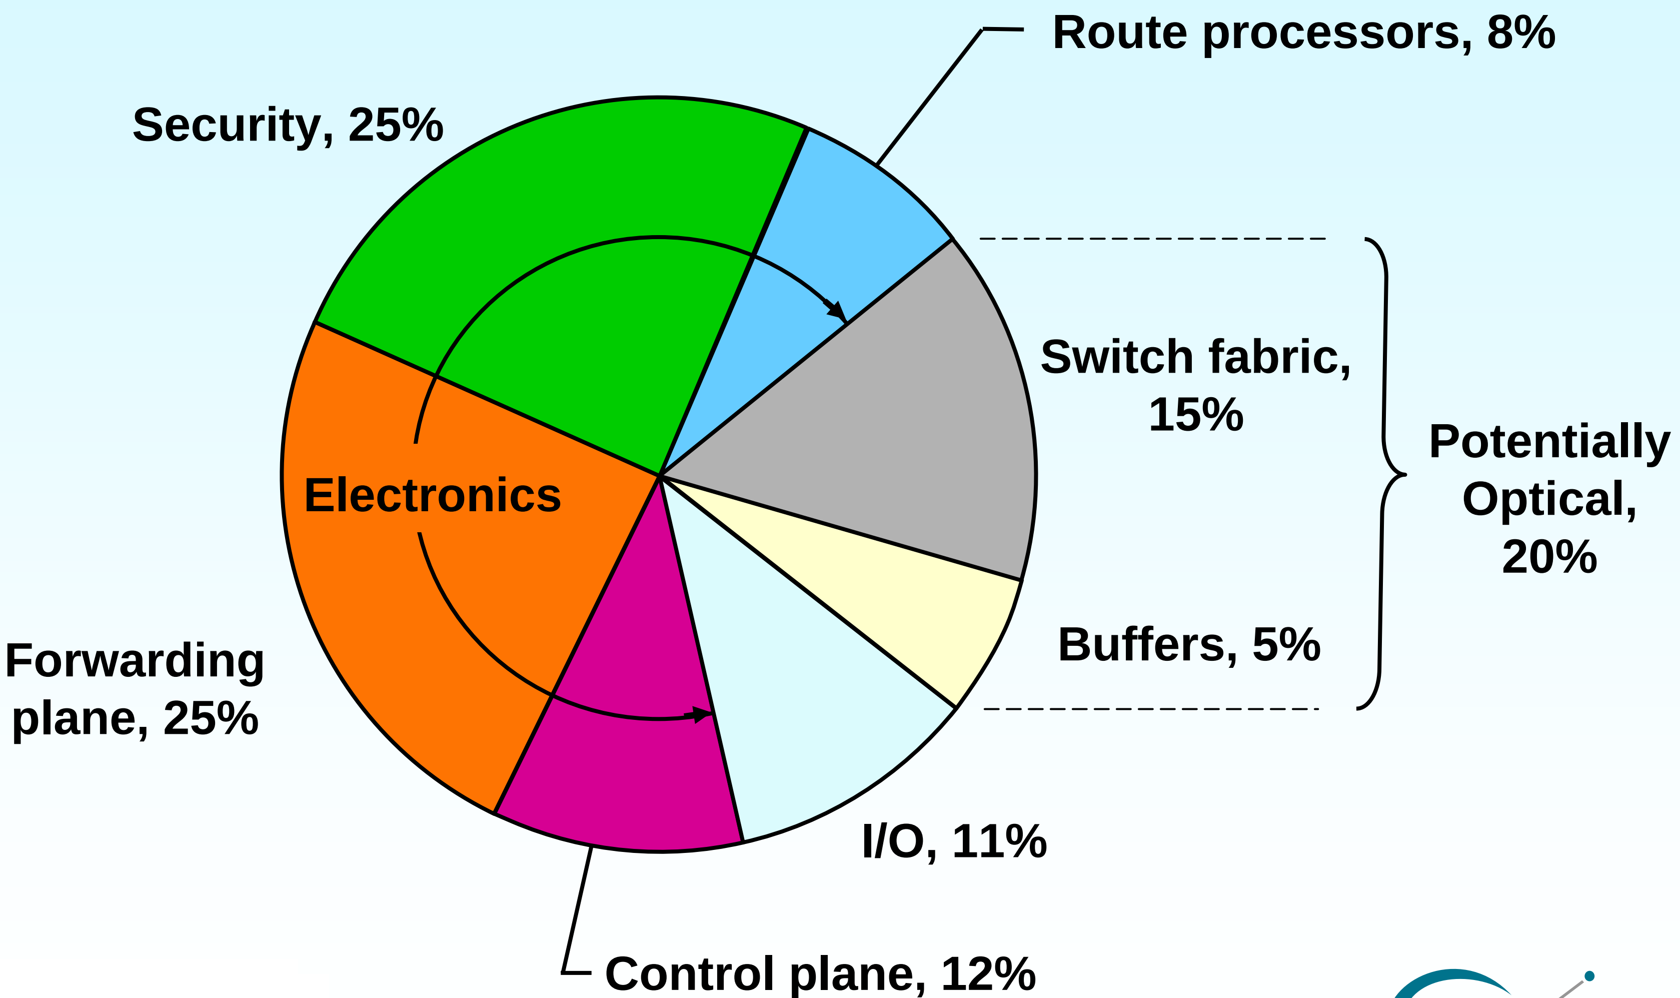

Source: G. Epps, Cisco

Router Architectures

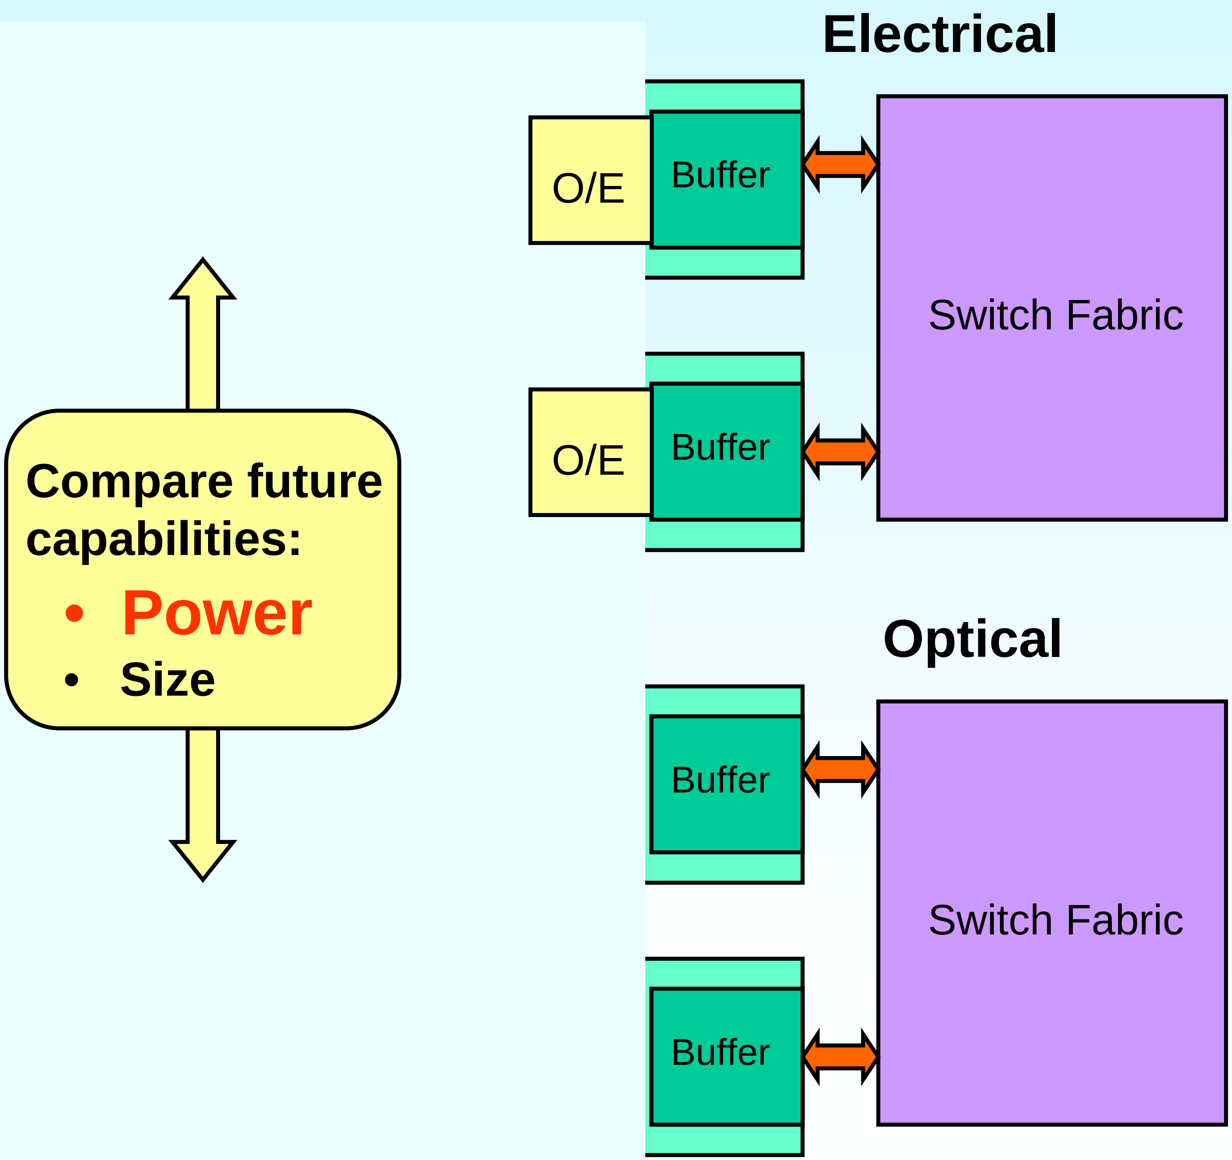

Methodology

"Making predictions is difficult....especially about the future." *Yogi Berra*

	Optics	Electronics
Buffers	▪ Fiber, Slow Light, Planar WG's ▪ Fundamental limitations -Storage Density, Power/Energy ▪ Waveguide attenuation	▪ 20-nm CMOS (~ 2020) International Technology Roadmap for Semiconductors http://public.itrs.net/
Switch Fabric	▪ SOA gate arrays ▪ AWG's and λ Converters (SOA) ▪ Performance projections - gain, size, attenuation ▪ Fundamental limitations - ASE, extinction ratio	

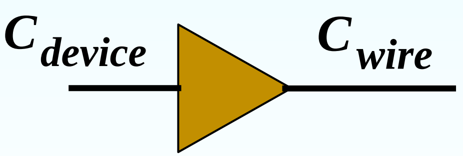

$$\text{Energy} = \frac{1}{2} C_{total} V^2$$

Loss Happens!

Loss → Energy

CUBIN

1 Pb/s Switch – "Traditional" Buffer Size

25,000 ports at 40 Gb/s, 250 ms buffering per port

EPS

CMOS Buffers

Total buffer capacity 250 Tb

~ 10^4 RAM chips @ 3 GB/chip

OPS

Fibre Delay Line Buffers

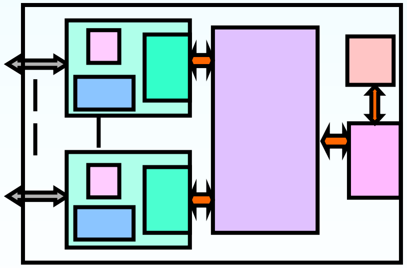

Total fibre length = 1 Tm

Distance from the Sun to Saturn

Packet Switching with Reduced Buffering

Enachescu et al., *ACM/SIGCOMM July 2005:* Buffer size can be reduced

OPS

Fibre Delay Line Buffers

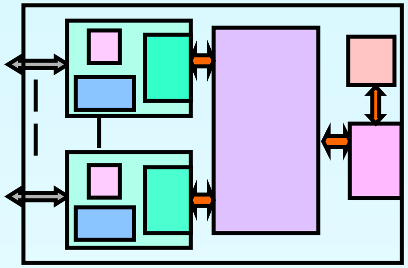

5 µs buffering per port (200 kb/port)
~20 packets @ 40 Gb/s

Total fibre length = 30 Mm (~1 km/port)

~ Circumference of the Earth

Buffering with fiber delay lines is a challenge

Slow Light to the Rescue?

Photonic Crystal

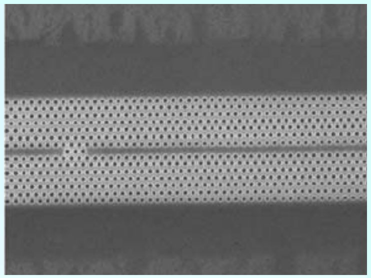

Group Velocity

$$v_g = \frac{c}{n + \omega \frac{dn}{d\omega}}$$

Large

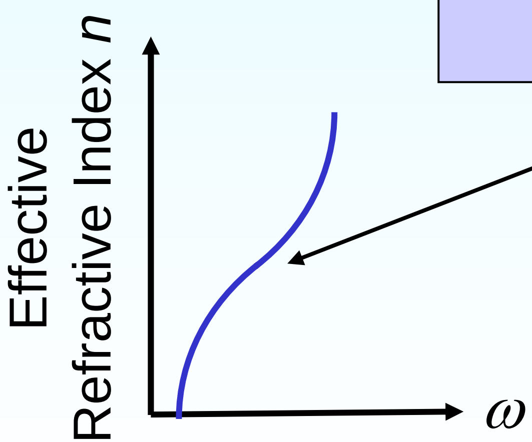

Y. Xu et al., *QELS*, 2000

Electromagnetically-Induced Transparency in Semiconductor

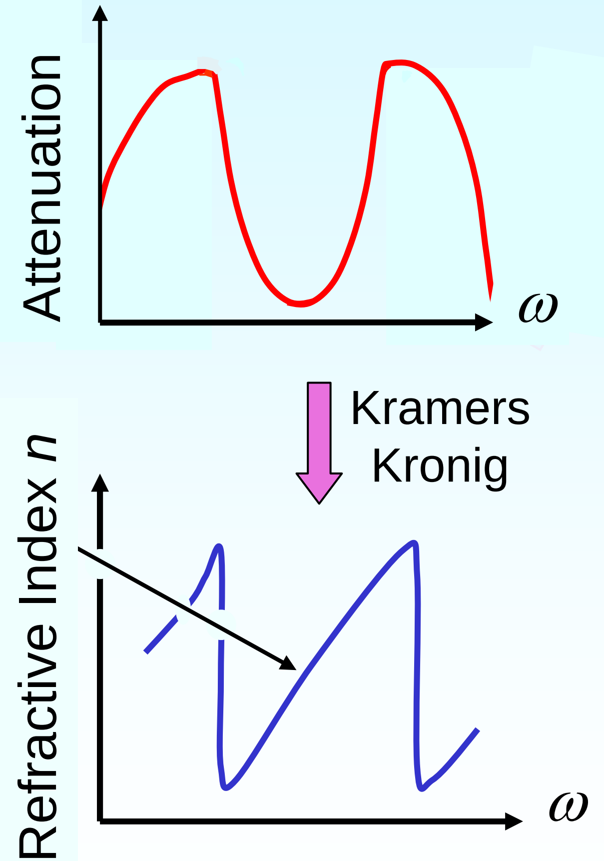

C.J Chang-Hasnain et al.,
Proc IEEE, **9**, 2003

Car Analogy

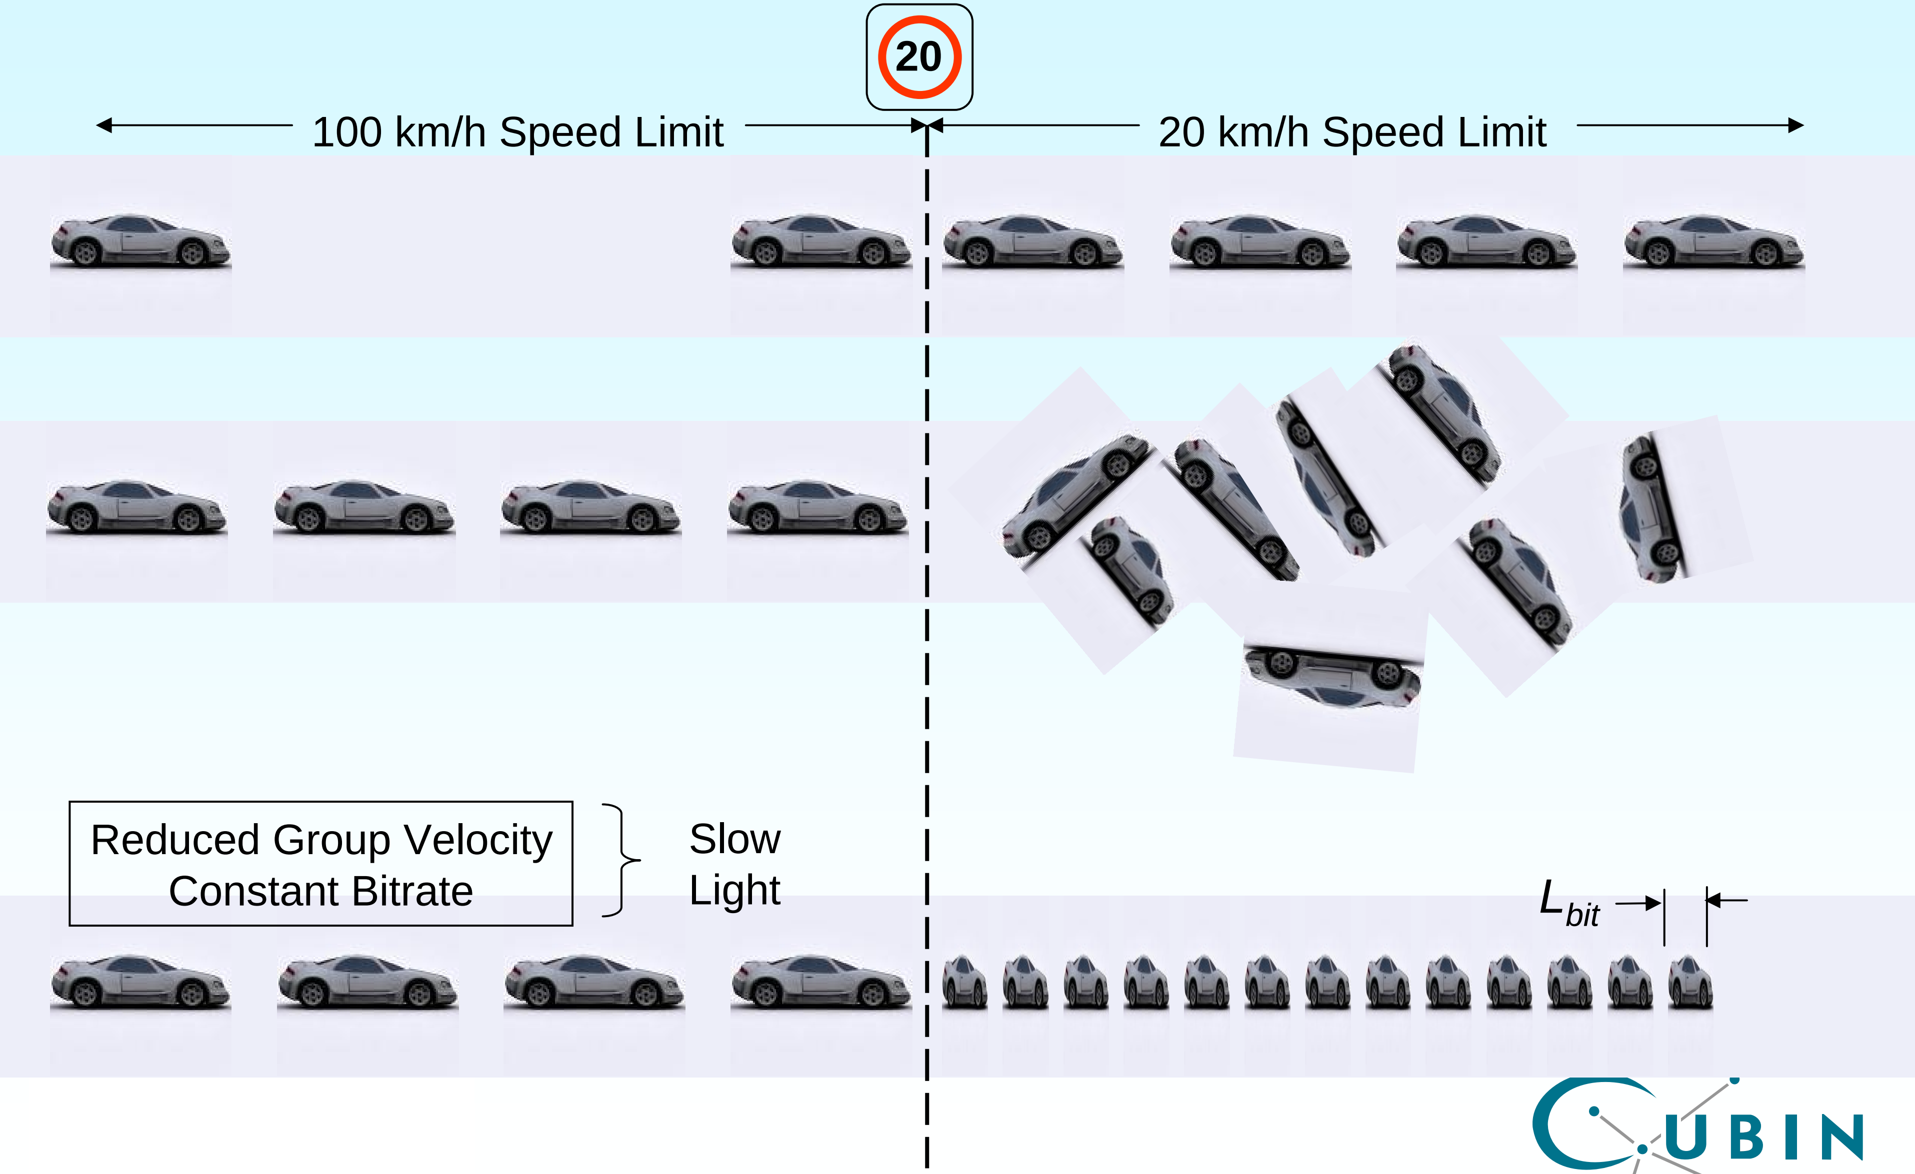
20
100 km/h Speed Limit
20 km/h Speed Limit
Reduced Group Velocity
Constant Bitrate
Slow
Light
L_{bit}

Some Numbers

100 -Tb/s Optical Router
1 ms buffering per port
(100 Gb total)

Fiber

200 km/port, 0.5 Gm total

Storage Density:
1 bit / 5 mm

Slow Light Waveguide

Slow-down factor = 5×10^2
400 m/port, 1 Mm total

Storage Density:
1 bit / 10 μm

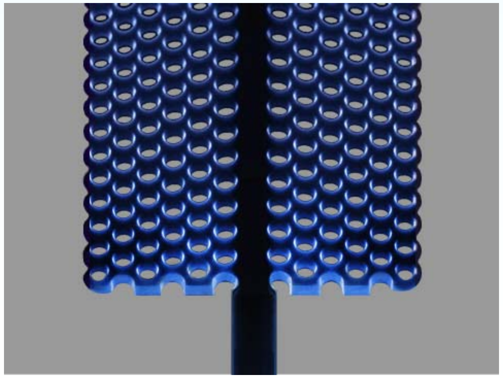

Ideal Slow Light Waveguide

40 m/port, 100 km total

Storage Density:
~1 bit / μm

Wavelength

CUBIN

Size Matters

Ideal Slow Light Waveguide

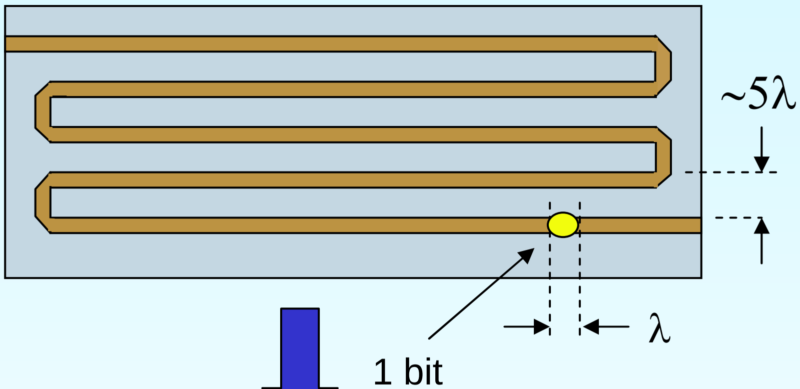

Minimum bit area ~ $5\lambda^2$
(λ = ~1 μm)

Storage Density: 150 Gbit/m²

CMOS (2018)

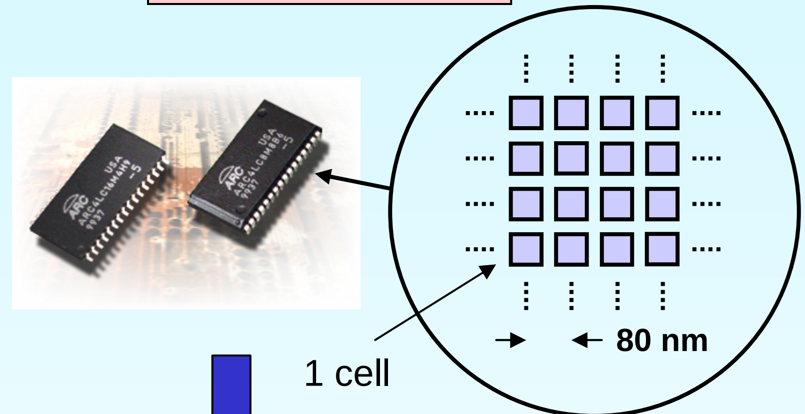

eDRAM cell area
80 nm x 80 nm

Storage Density: 150 Tbit/m²

per wavelength

Area	Capacity	Area
700 cm²	← 100 Gbit →	0.7 cm²
17 m²	← 2.5 Tbit →	170 cm²

Loss Happens!

Fibre: ~0.2 dB/km

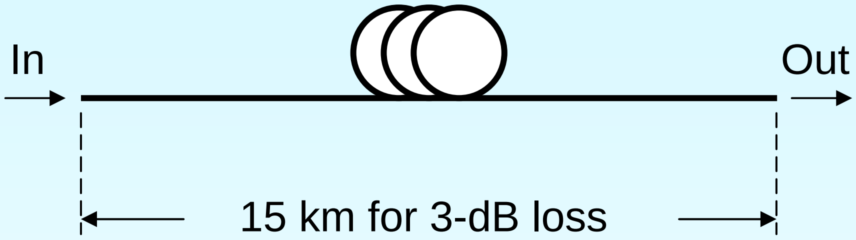

Planar WG: 0.1 dB/cm

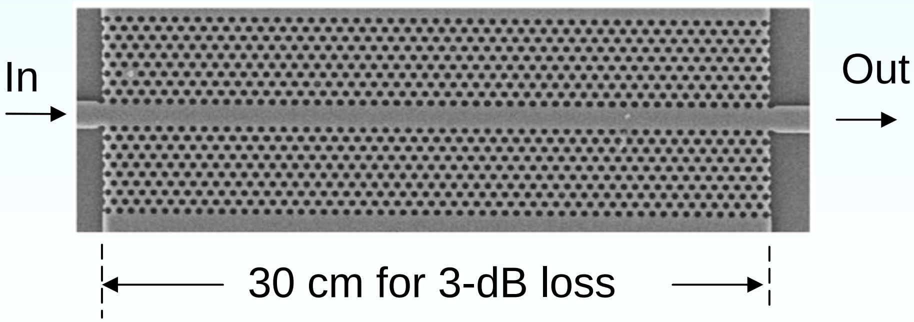

200 m (200 packets) ⇨ 33,000 dB

20 m (20 packets) ⇨ 3,000 dB

QUBIN

Dispersion

Photonic Crystal

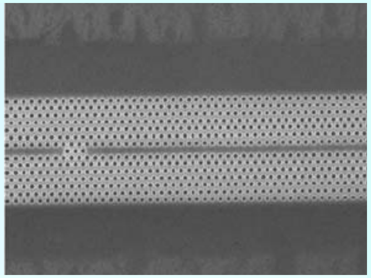

Electromagnetically-Induced Transparency in Semiconductor

Group Velocity

$$v_g = \frac{c}{n + \omega \frac{dn}{d\omega}}$$

Changes across passband

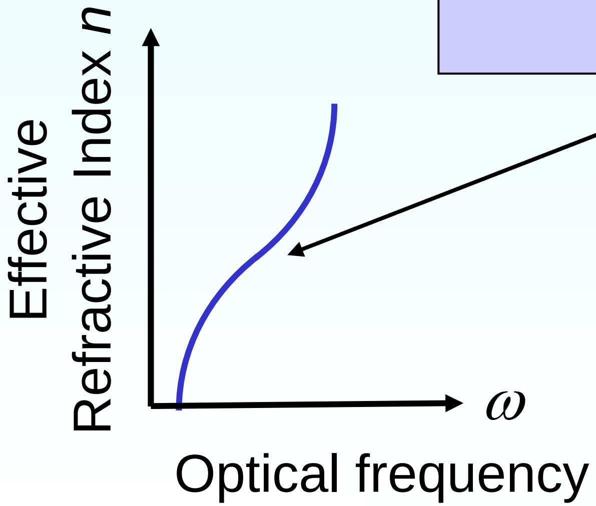

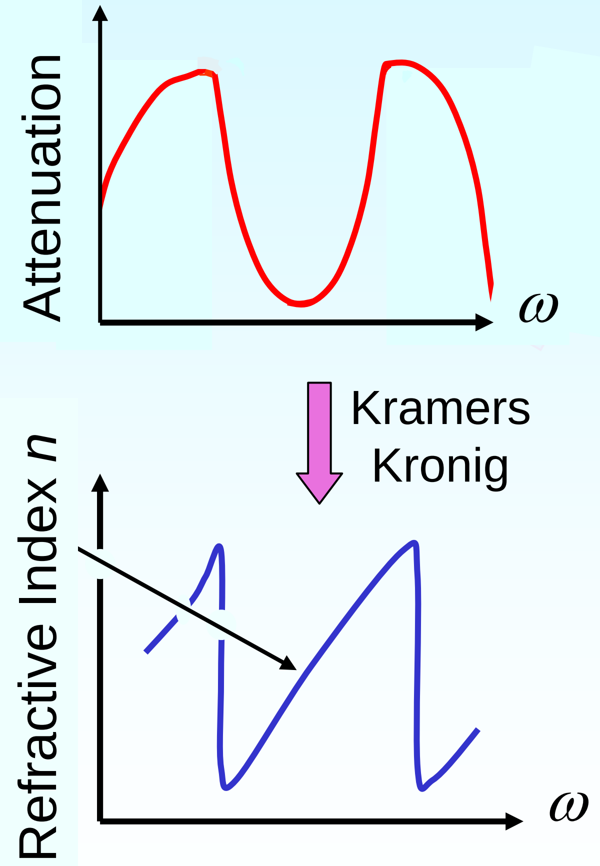

Y. Xu et al., *QELS*, 2000

C.J Chang-Hasnain et al., *Proc IEEE*, **9**, 2003

CUBIN

Putting it All Together: Delay Line Buffer

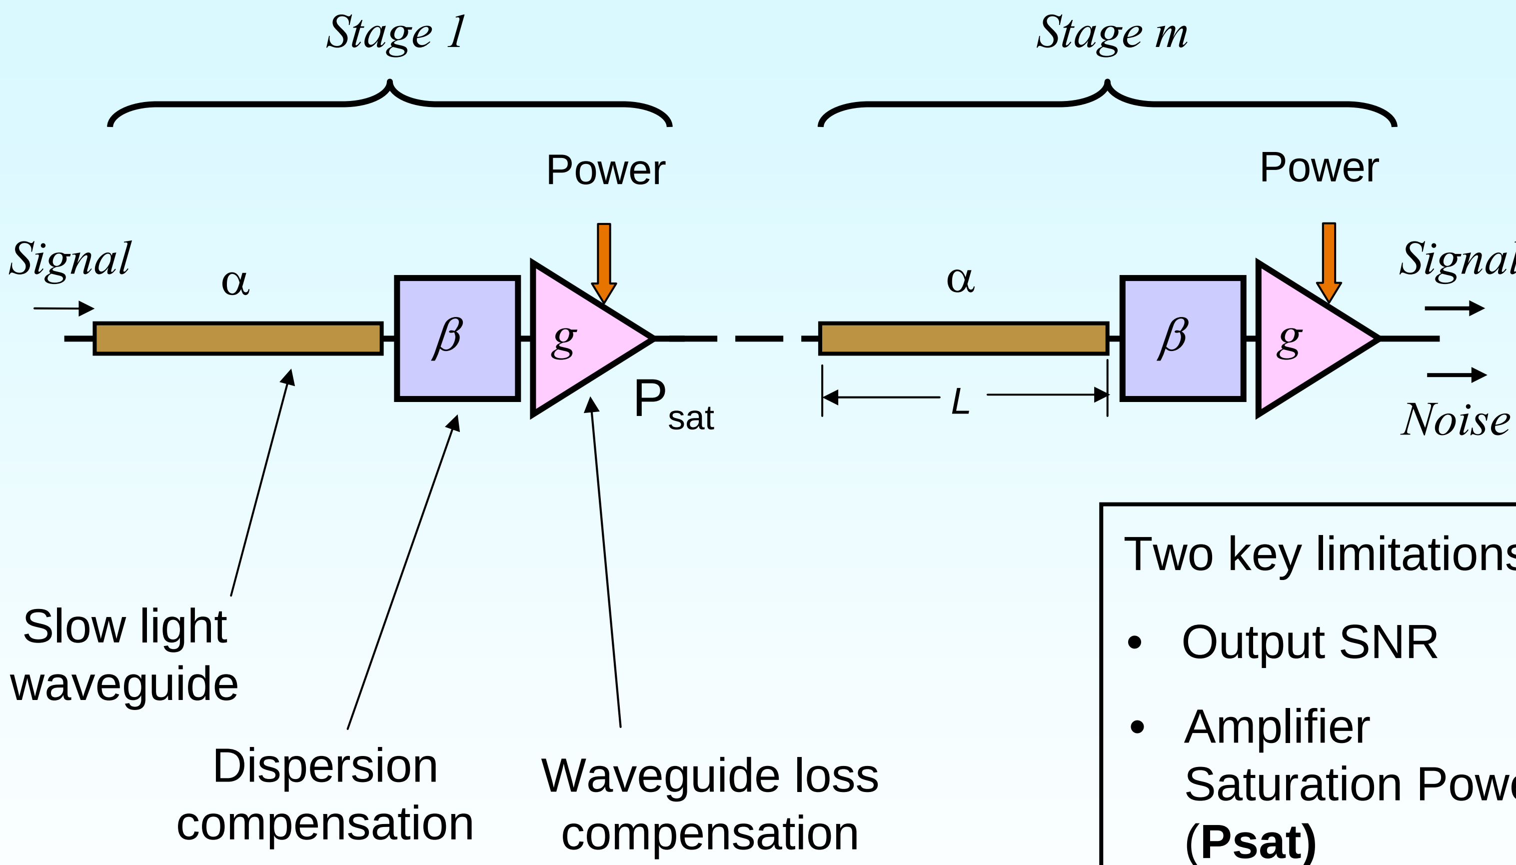

Two key limitations:

- Output SNR
- Amplifier Saturation Power **(Psat)**

Buffer Capacity

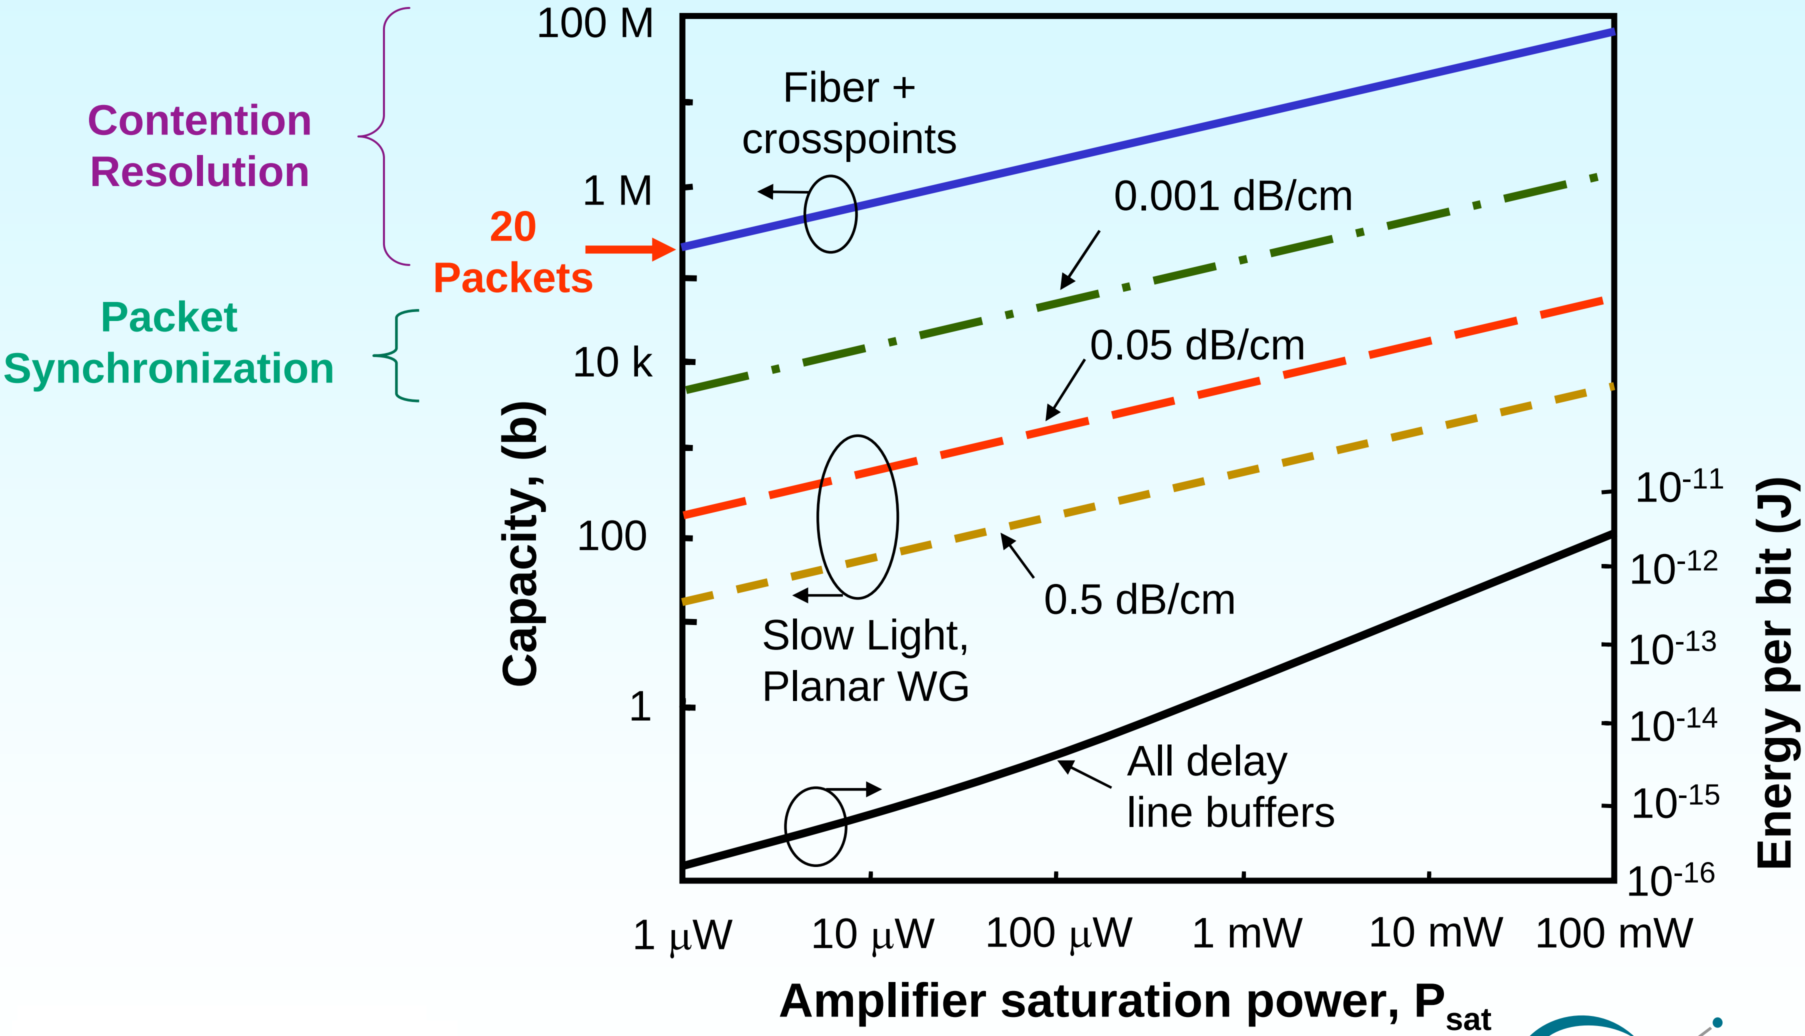

Methodology

	Optics	Electronics
Buffers	▪ Fiber, Slow Light ▪ Fundamental limitations -Storage Density, Power/Energy ▪ Waveguide attenuation	▪ 20-nm CMOS (~ 2020) International Technology Roadmap for Semiconductors http://public.itrs.net/ More conservative projections than for Optics
Switch Fabric	▪ SOA gate arrays ▪ AWG's and λ Converters (SOA) ▪ Performance projections - gain, size ▪ Fundamental limitations - ASE, extinction ratio	

QUBIN

Energy per Bit in Optical Switches

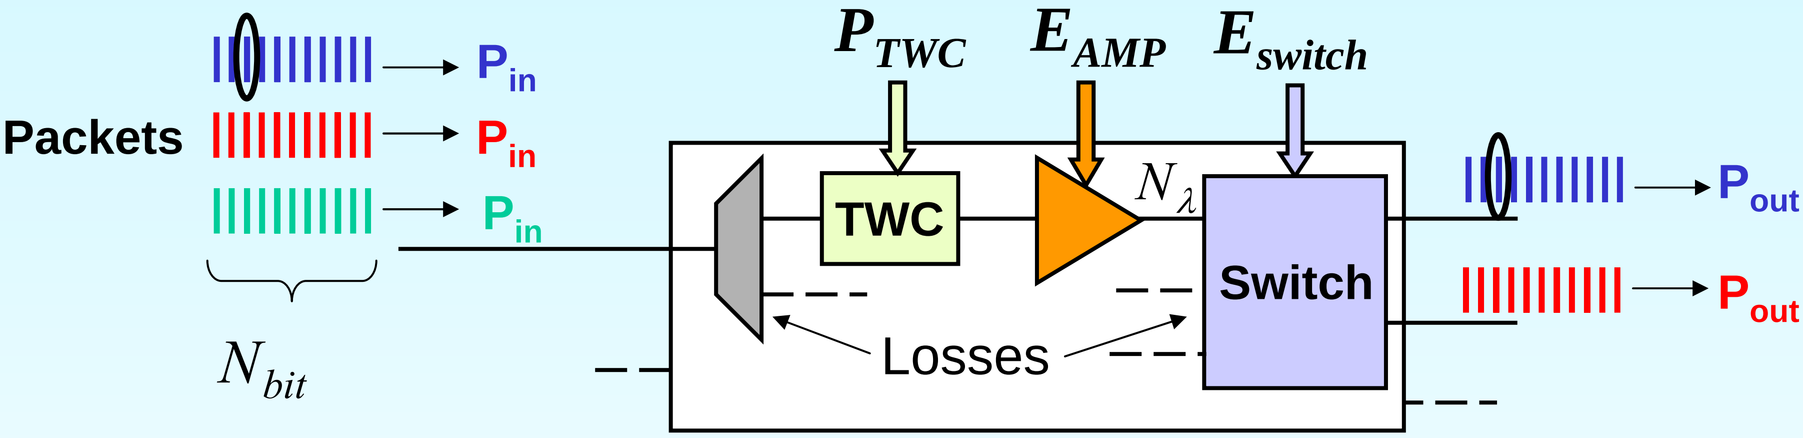

Energy per bit

↓ as N_λ and N_{bit} ↑

Losses

Large in SOA gate arrays

$$E_{bit} = \frac{E_{switch}}{N_\lambda N_{bit}} + (P_{in} - P_{out})\tau_{bit} + E_{TWC} + E_{AMP}$$

Conventional “wisdom”

Number of λ's

Number of bits

Dominant terms in large fabrics

1 Pb/s Cross-Connect (40 Gb/s)

3-stage CLOS

AWG-Based **8 pJ/bit**

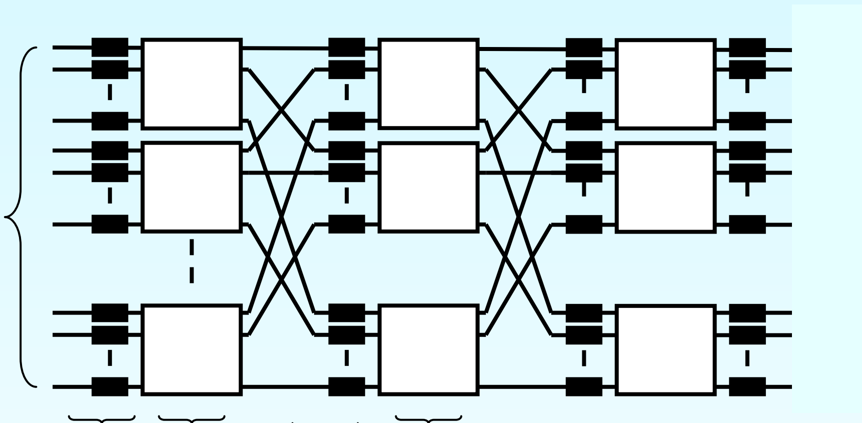

CMOS **14 pJ/bit**

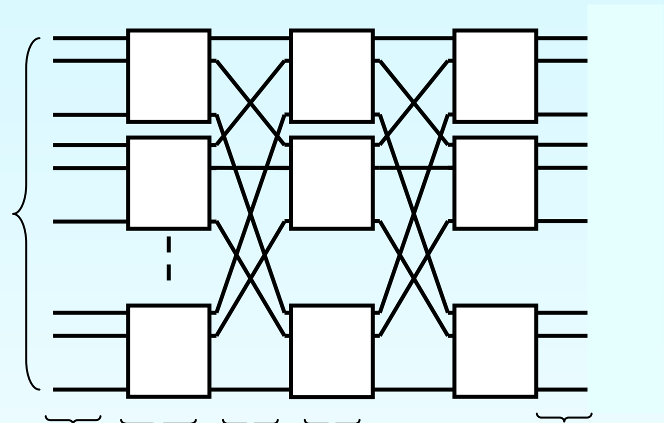

SOA-Based **>20 pJ/bit**

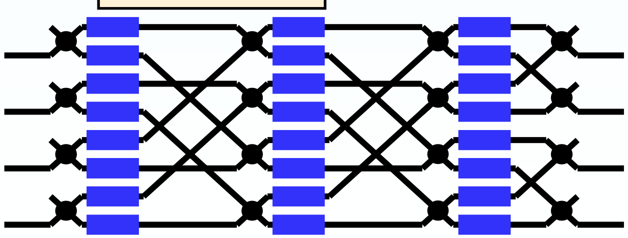

P_{sat} > 100 mW

Building Block:
Benes Switch Array

Kalman et al., *PTL*, **4**, 1992

CUBIN

Comparing Cross-Connects : 1-Pb/s Router

	Power Dissipation		Device Area (not including TWC's and interconnects)	
Switch Rate	40 Gbit/s	160 Gbit/s	40 Gbit/s	160 Gbit/s
SOA Gate Array	> 10 kW	> 20 kW	80 m²	5 m²
AWG-Based Switch	2.3 kW	4.6 kW	>1 m²	<1 m²
CMOS	4.6 kW	-	0.2 m²	-

- AWG has the "edge (using OEO WC's)
- Dominates over Buffer

CMOS has the "edge"

Buffer/Switch Combinations

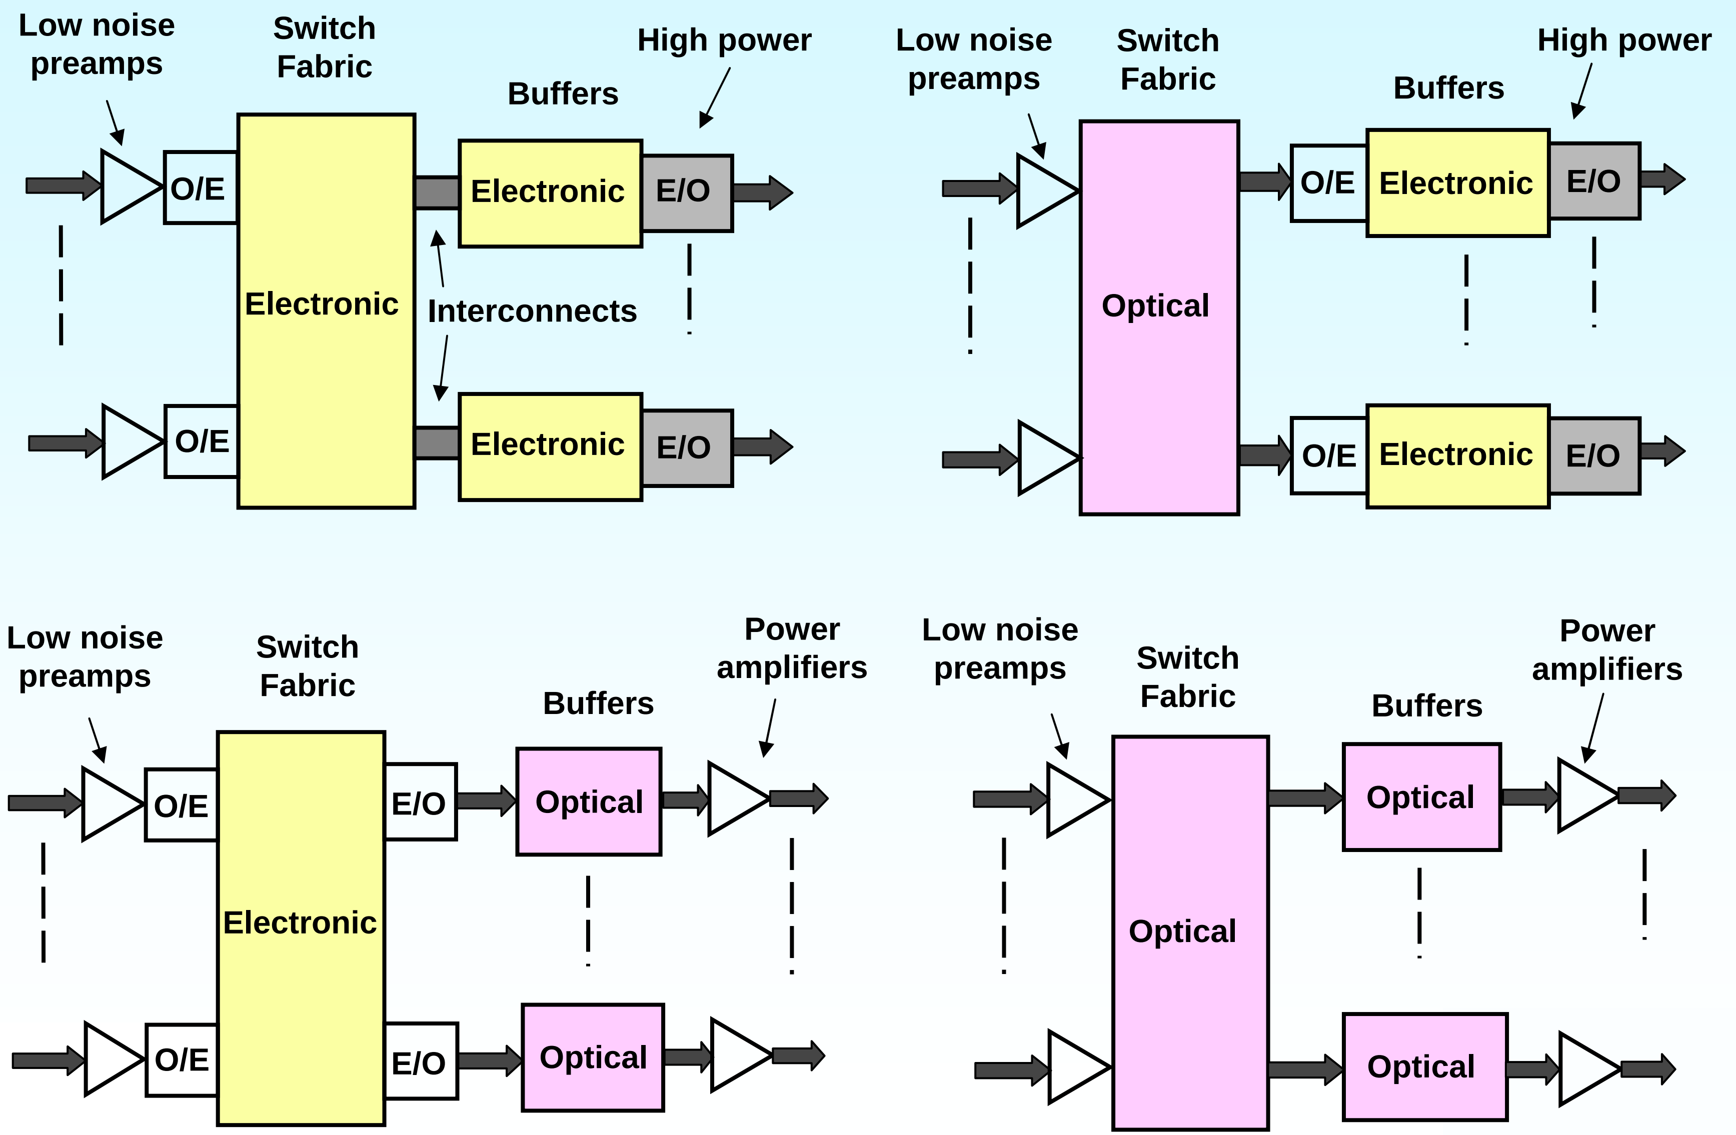

1-Pb/s Router Power and Area

Buffer: 400 kb per port

			Buffer							
			Optical				Electronic			
			Buffer	Cross connect	Extra power	Total	Buffer	Cross Connect	Extra power	Total
Cross connect	Optical	Power	2 kW	2.3 kW	2.6 kW	6.9 kW	0.2 kW	2.3 kW	2.5 kW	5.0 kW
		Area	125 m^2	1 m^2	-	126 m^2	50 cm^2	1 m^2	-	1 m^2
	Electronic	Power	2 kW	4.6 kW	2.2 kW	8.8 kW	0.2 kW	4.6 kW	2.1 kW	6.9 kW
		Area	125 m^2	1 m^2	-	126 m^2	50 cm^2	1 m^2	-	1 m^2

Smallest area

Lowest Power

Conclusion:

There is no compelling case for all-optical routers

Conclusions

- Petabit-per-second optical routers
  - will become technically feasible
  - optoelectronic rather than all-optical

- Power consumption is *the* central design issue in large routers
  - forwarding, control, and security processing
  - buffers
  - switch fabric
  - interconnects

  } Energy per bit is a key parameter (buffers, switch fabric, interconnects)

- Optical solutions require very low loss waveguides (< 0.1 dB/m)

- Is the "electronic bottleneck" a myth?

CUBIN