Reality and Fantasy in Optics

ECE 594 Special Topics in Electrical and Computer Engineering
University of California Santa Barbara
Guest Lecture

29 May 2026
Chris Cole

Outline

- ➢ **Optics & AI**
- ▪ Optics Modulation Alternatives
- ▪ Optical Computing (ref: https://opg.optica.org/oe/fulltext.cfm?uri=oe-29-9-13153)
  - Frenzy Examples
  - Historical Examples
  - Precision
  - Comparison
  - Addition
  - Multiplication
  - Convolution
  - Discussion
- ▪ Optics: Reality or Fantasy

Optics & AI

- Next Datacom Paradigm shift:
  - Optical Computer I/O driven by AI/ML
  - Hottest segment of the AI investment bubble
- Optical Computer I/O requirements:
  - Order(s) of magnitude more stringent than Optical Networking
  - Only met with fundamentally new optical Components and Devices
- Datacom Optics investment unfortunate concentration:
  - Sub-systems and Systems
  - Rearranging and/or aggregating existing optical technology
- Real optics innovation is expensive, long term, and process based
  - Most investors do not like to hear this

Datacom Paradigm Shift Enabling Optical Technologies

Datacenter Paradigm	Network or Computer Link Rate	No. of Lanes	Enabling Component & Device Technology	Enabling Sub-system Technology
Enterprise	100M (ex. Ethernet) 1G 10G	1	VCSEL DFB LASER	LC (Lucent Connector) Pluggable Module
Hyperscale	40G 25/50G 100G	4	EML WDM Si MZM	MT Parallel Connector
	200G 400G 800G	4, 8	DSP	Heatsink 😀
AI/ML	1.6/3.2T 6.4T 12.8T	≥ 8	Hi-Rel LASER DWDM High Density Modulator	LSI SiPh LSI Optics Packaging High Density Connector

Outline

- Optics & AI
- ➢ **Optics Modulation Alternatives**
- Optical Computing
  - Frenzy Examples
  - Historical Examples
  - Precision
  - Comparison
  - Addition
  - Multiplication
  - Convolution
  - Discussion
- Optics: Reality or Fantasy

Datacom Optics Modulation Alternatives: NRZ & PAM4

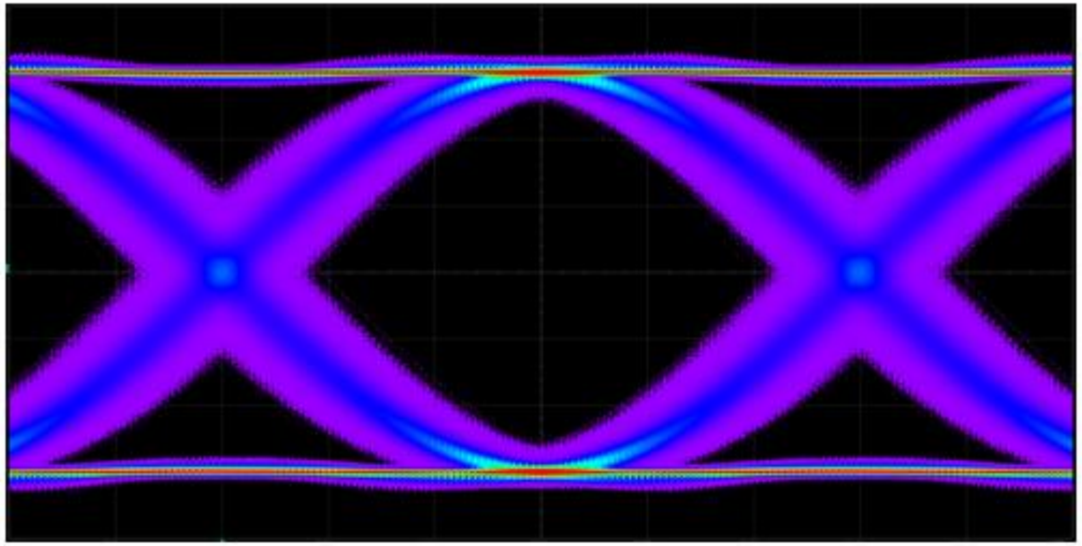

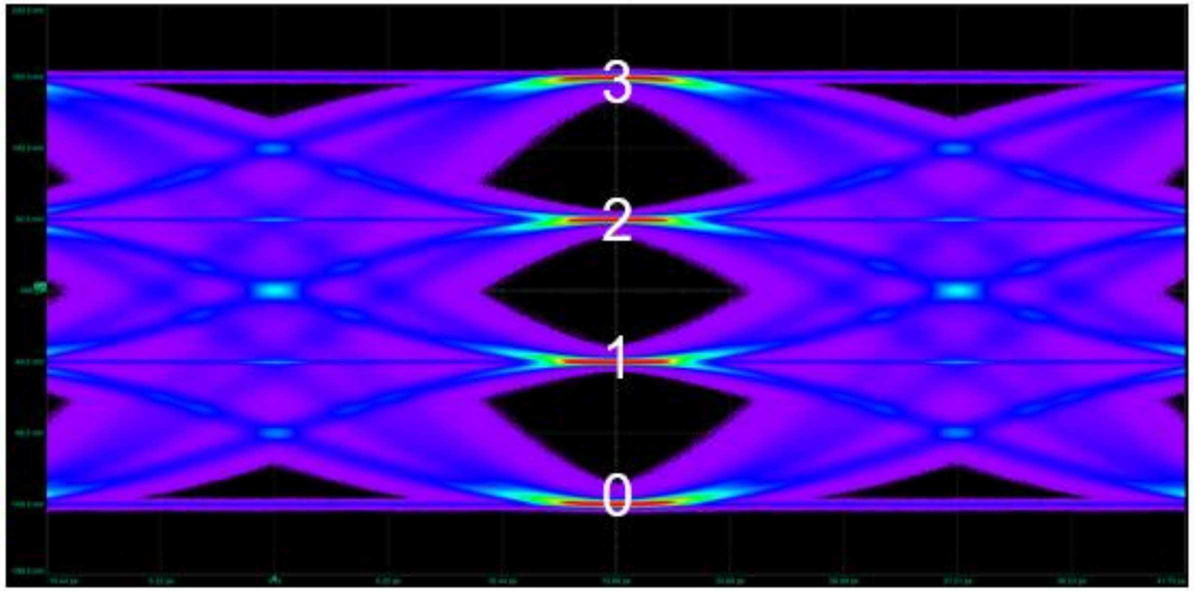

Datacom Optics NRZ vs. PAM4 Debate

- NRZ was used on all 25G and lower speed λs
  - 1G λs : 1GbE
  - 10G λs: 10GbE, 40GbE
  - 25G λs: 25GbE, 100GbE
- IEEE chose PAM4 modulation for 50G copper lanes in 2012
- NRZ vs. PAM4 50G λs debate started in IEEE in 2012
  - PAM4 25Gbaud adopted in 2015 in 200GbE & 400GbE project
  - Enabled reuse of 50G PAM4 SerDes technology in development for ASICs
  - Enabled reuse of 25Gbaud tech for perceived quicker time to market
- PAM4 has been used for $\geq$ 50G λs (50, 100, 200, 400G)

Shannon-Hartley Theorem

- $C = B \log_2(1 + S/N)$
  - $C \triangleq$ Channel capacity
  - $B \triangleq$ Bandwidth
  - $S \triangleq$ Signal Power
  - $N \triangleq$ Noise Power
- Guidance to increase C:
  - If B limited, increase S/N to support higher order modulation (HOM)
  - If S/N limited, increase B to support higher Baud rate

Channel Loss & TRX S/N

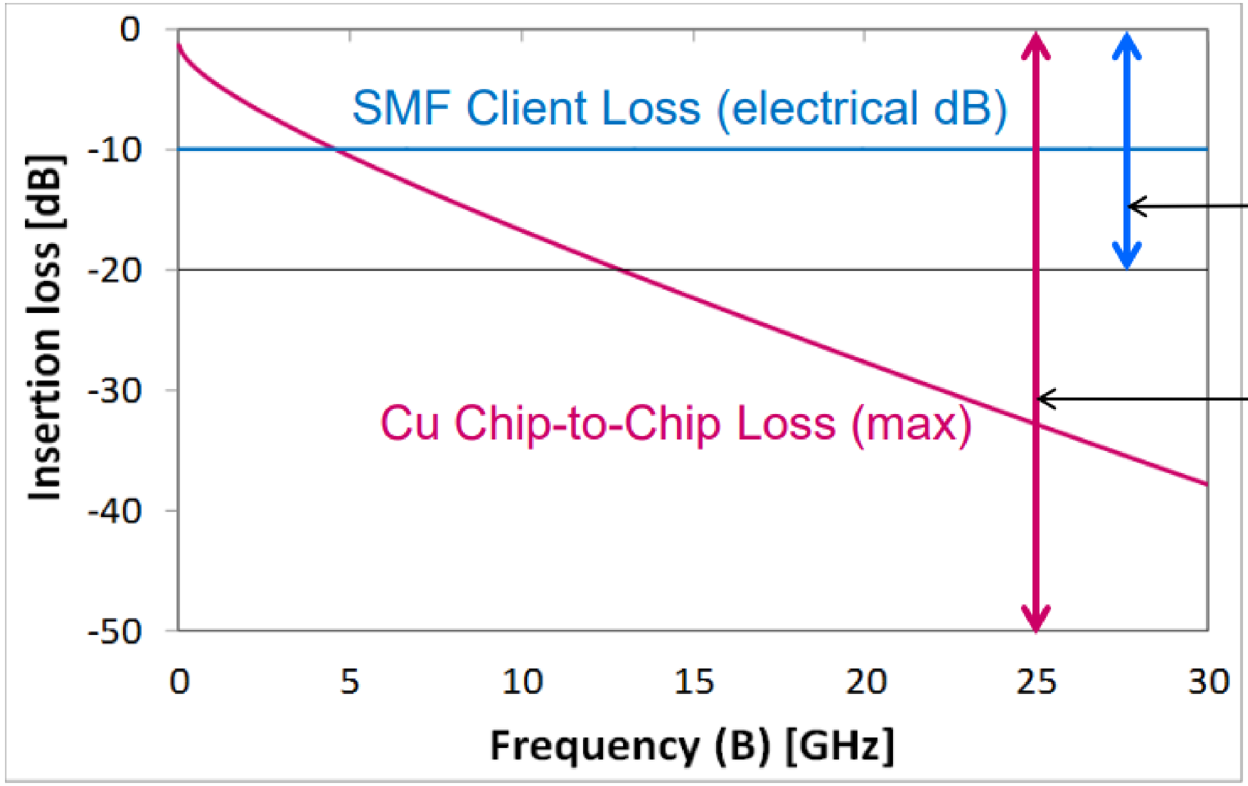

S/N no FEC BTB	Limitation		Modulation Guide
	Channel B	TRX S/N	
SMF TRX	No	Yes	NRZ
Cu SerDes	Yes	No	HOM

Component Bandwidth & Vertical Eye Closure (VEC)

VEC improves with component bandwidth (B) which increases over time.

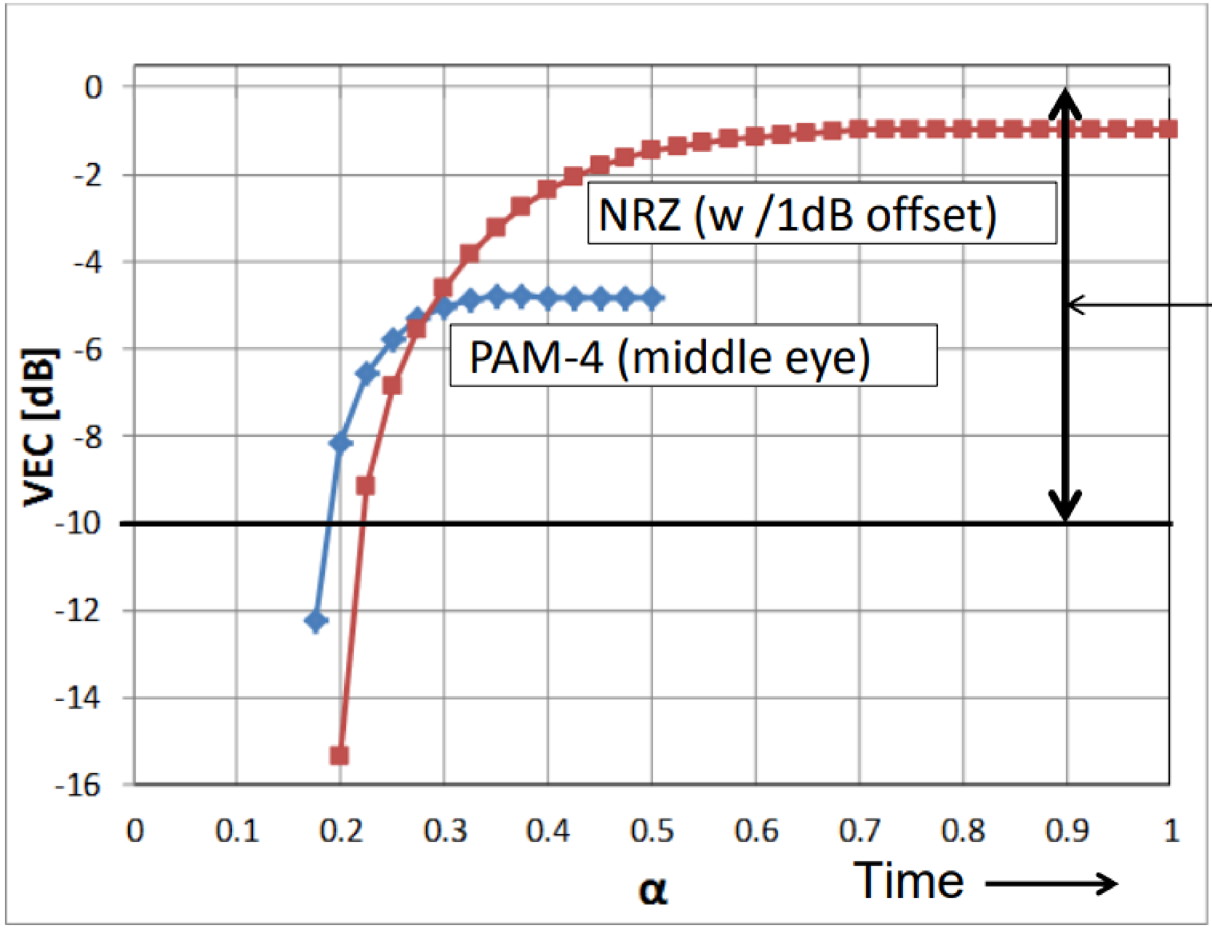

noise penalty offsets VEC by ~1dB (B_{NRZ}/B_{PAM-4} dependent)

α = B / bit-rate

Discussion

- For datacom optics, NRZ is the preferred choice if feasible, because it has the highest margin
- If single λ NRZ insufficient, Parallel and/or WDM are the preferred to increase rate
- As component bandwidth increases with time, NRZ optics margin improves the most which drives down cost

  (ex. 10G Serial NRZ optics)
- PAM-4 permanently locks in ~3dB S/N penalty limiting optics margin improvement, even as component bandwidth increases

Did the Optics Industry Blunder by Switching from NRZ to PAM4?

- It sure did!
- 25GBaud PAM4 reduced the cost & time to market of initial low-volume shipments
- Predictably optical component bandwidth increased over time
- 50GBaud technology matured and is now shipping for 100G PAM4
- 25Gbaud PAM4 λ optics shipped millions of ports
- 50Gbaud NRZ λ optics, if adopted, would instead have shipped millions of ports
- PAM4 λ optics 3dB S/N penalty permanently locked-in
- Significant power and cost penalty for 50G, 100G and 200G λ optics

Outline

- Optics & AI
- Optics Modulation Alternatives
- Optical Computing (ref: https://opg.optica.org/oe/fulltext.cfm?uri=oe-29-9-13153)
    - ➢**Frenzy Examples**
    - Historical Examples
    - Precision
    - Comparison
    - Addition
    - Multiplication
    - Convolution
    - Discussion
- Optics: Reality or Fantasy

Optical Computing Ex. 1 & 2

- Frequent pitch slogan: “Computing at the speed of light” (unlike electrons ☺)
- Lightmatter, the leader in **photonic supercomputing**, announced today it has raised a $400 million Series D, valuing the company at $4.4 billion and bringing the total capital raised to date to $850 million. The round was led by new investors advised by T. Rowe Price, with participation from existing investors, including Fidelity
- Lightelligence, a Shanghai-based specialist in optical interconnect and **optical computing**, delivered an explosive performance during their Hong Kong stock market debut, underscoring strong investor enthusiasm for artificial intelligence (AI) and next-generation semiconductor technologies. The company’s stock surged by as much as 408%, reaching HK$930 compared to its initial offering price of HK$183.2, according to data from the Hong Kong Exchange.

Optical Computing Ex. 3

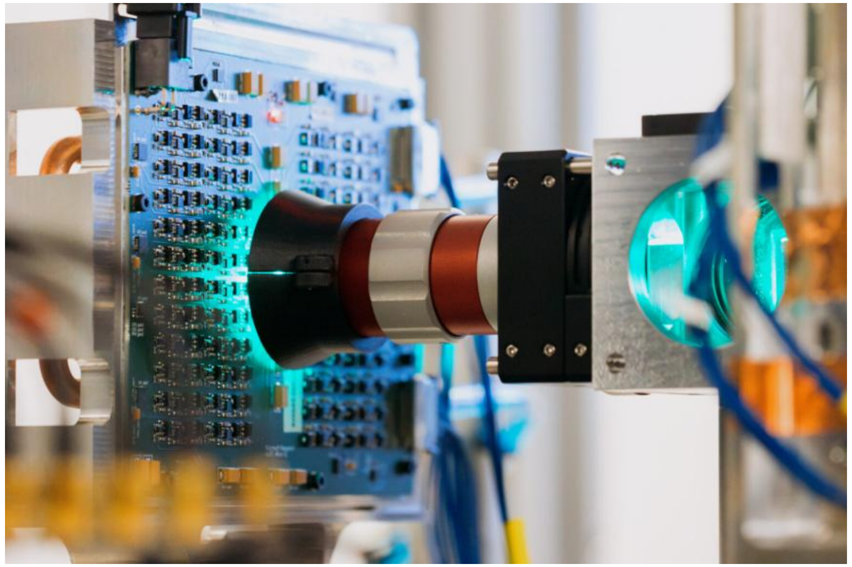

- Satya Nadella, Microsoft CEO LinkedIn post:

  “Our breakthrough work on an analog **optical computer** points to new ways to solve complex real-world problems with much great efficiency.”

- Chris Cole LinkedIn response comment:

  “Practical is the operative word:

  a massive 3D fine tuned set-up, ~billion times bigger than an equivalent functionality CMOS chip.”

- MS industry colleague private comment:

  “MS is not likely any time soon to ask for your consulting services.” ☺

Optical Computing Ex. 4

- Neurophos, a leader in **photonic AI chip** technology, has raised $110 million in an oversubscribed Series A round, bringing total funding to $118 million. The round was led by Gates Frontier, with participation from M12 (Microsoft’s Venture Fund), … and others
- “Modern AI inference demands monumental amounts of power and compute,” said Dr. Marc Tremblay, Corporate Vice President and Technical Fellow of Core AI Infrastructure at Microsoft. “We need a breakthrough in compute on par with the leaps we've seen in AI models themselves, which is … Neurophos' technology.”

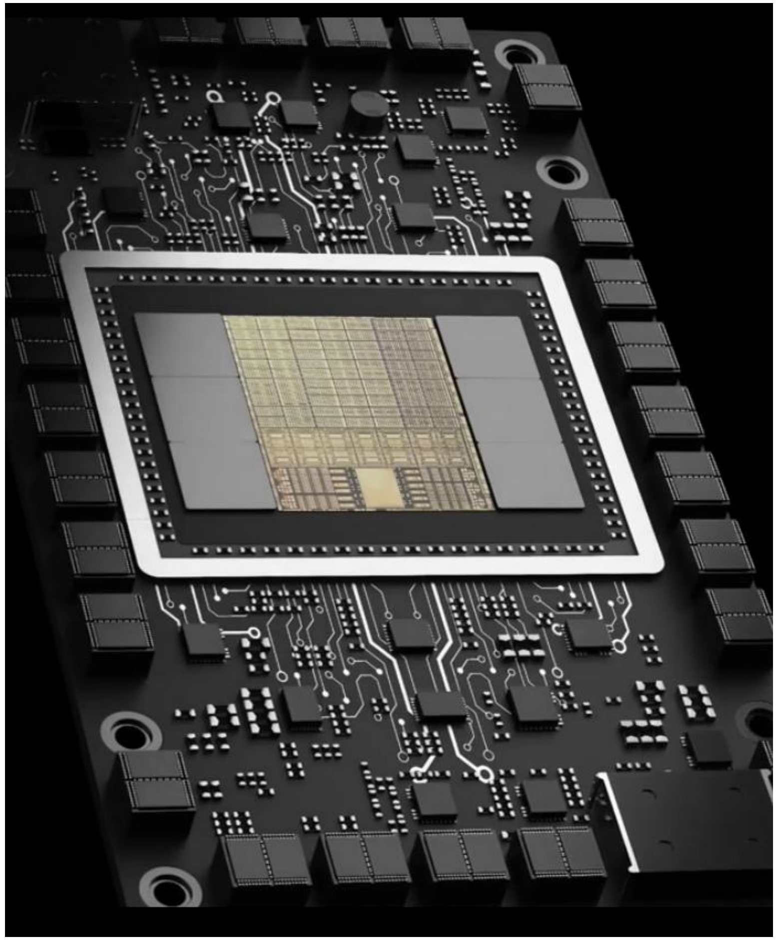

Outline

- Optics & AI
- Optics Modulation Alternatives
- Optical Computing
  - Frenzy Examples
  - ➢**Historical Examples**
  - Precision
  - Comparison
  - Addition
  - Multiplication
  - Convolution
  - Discussion
- Optics: Reality or Fantasy

1^{st} Optical Computing Example: 4600-year-old Egyptian Lens

- The “eyes” appear to follow the observer as they move about the statue
- On display at Louvre Museum, Paris

Widely Used Optical Computing Example: Eyeglasses

- Two lenses in a wooden frame, Italy, 1280’s
- Lens processing is 2-D spatial filtering or 2-D convolution, i.e., inference
- A hypothetical electronic lens processes 24-bit RGB 512x512 pixel image at 120 frames/sec
  - → ~25 trillion 8-bit Multiply-Accumulates/sec
- Zero incremental energy

Widely Used Optical Computing Example: Eyeglasses

- Two lenses in a wooden frame, Italy, 1280's
- Lens processing is 2-D spatial filtering or 2-D convolution, i.e., inference
- A hypothetical electronic lens processes 24-bit RGB 512x512 pixel image at 120 frames/sec
  - → ~25 trillion 8-bit Multiply-Accumulates/sec
- Zero incremental energy

- Problem: fixed focus (fixed coefficients)
- Solution: Ben Franklin bi-focal eyeglasses
  - → 1 bit of programmability

Telecom Optical Computing Example: DCF

- DCF (Dispersion Compensation Fiber) used in every Telecom link in the ’90s
- Passive, complex optical signal processing (computing)
- Zero incremental energy use (ignoring amplification for loss)
- Fixed compensation; requires a custom length spool for every link
- Dominant compensation approach despite extensive R&D into alternatives

Telecom Optical Computing Example: DCF

- DCF (Dispersion Compensation Fiber) used in every Telecom link in the ’90s
- Passive, complex optical signal processing (computing)
- Zero incremental energy use (ignoring amplification for loss)
- Fixed compensation; requires a custom length spool for every link
- Dominant compensation approach despite extensive R&D into alternatives
- Coherent DSP CMOS ASIC with adaptive equalization introduced 20 years ago
- Over time, completely replaced DCF and all other optical compensation techniques

Outline

- Optics & AI
- Optics Modulation Alternatives
- Optical Computing
  - Frenzy Examples
  - Historical Examples
  - ➢**Precision**
  - Comparison
  - Addition
  - Multiplication
  - Convolution
  - Discussion
- Optics: Reality or Fantasy

Generic Model of Mixed Signal Computing System

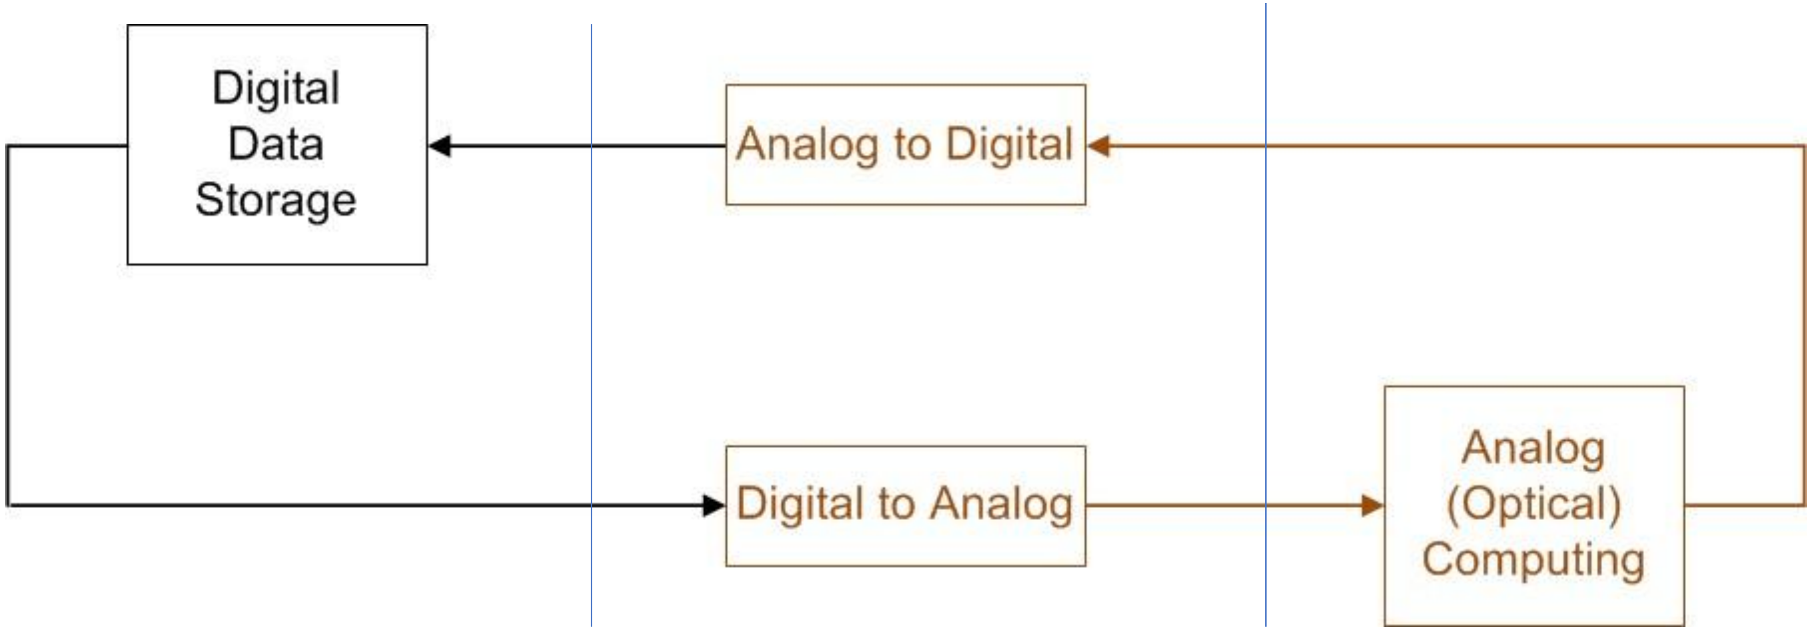

- ENOB (effective number of bits) is a fundamental figure of merit of mixed systems
- ENOB has been measured for decades:
  - SNR calculated from FFT
  - ENOB = $(\mathrm{SNR}_{\mathrm{dB}} – 1.76)/6.02$

Optical Compute HW

- Realistic Precision: ~4 bits,
  - INT4 not useful standalone in the Datacenter
- Claims of higher precision, like 8 bits, unsubstantiated
  - Rigorous measurement of ENOB never reported
- A priori, the required precision of AI applications is unknown and unpredictable
- Digital CMOS HW is programmable INTm and FPn, low-to-high precision
  - If useful, INT4 or INT2 only CMOS xPU would be tiny and trivial do develop
  - Same for other architectures like systolic array or in-memory HW computing

Outline

- Optics & AI
- Optics Modulation Alternatives
- Optical Computing
  - Frenzy Examples
  - Historical Examples
  - Precision
  - ➢**Comparison**
  - Addition
  - Multiplication
  - Convolution
  - Discussion
- Optics: Reality or Fantasy

Data Transfer Compute Model Optimized for Math Operations

Electrical Computing w/ electrical DT

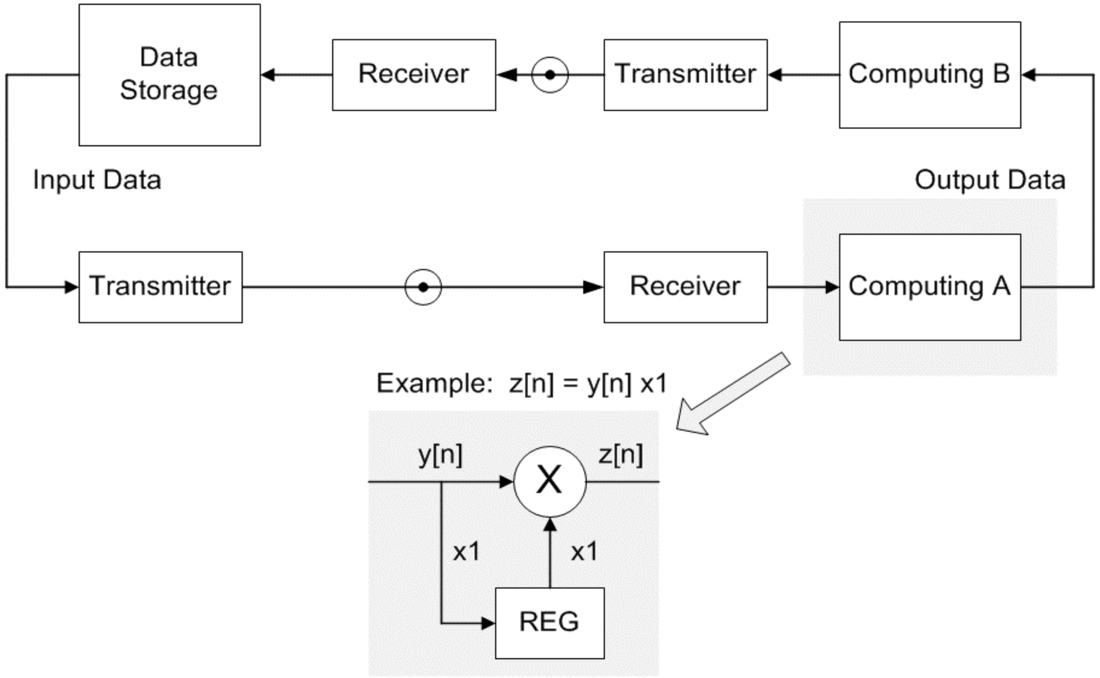

DT: Data Transfer

black: electrical elements

Apples-to-Oranges Energy Use Comparison Models

Electrical Computing w/ electrical DT

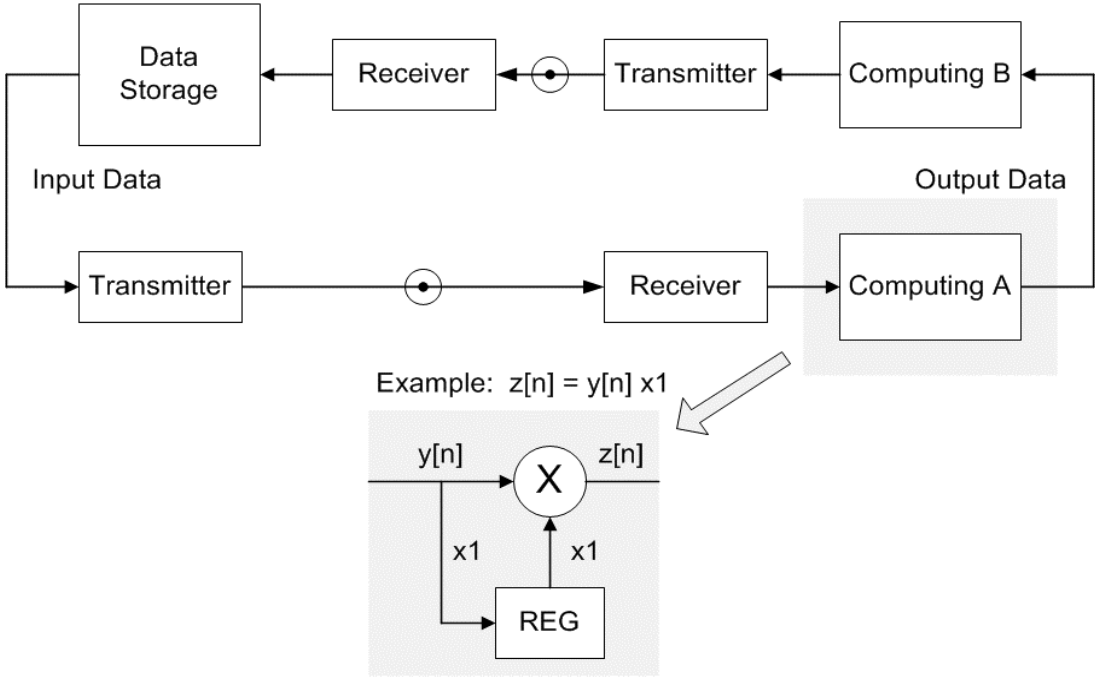

Optical Computing w/ optical DT

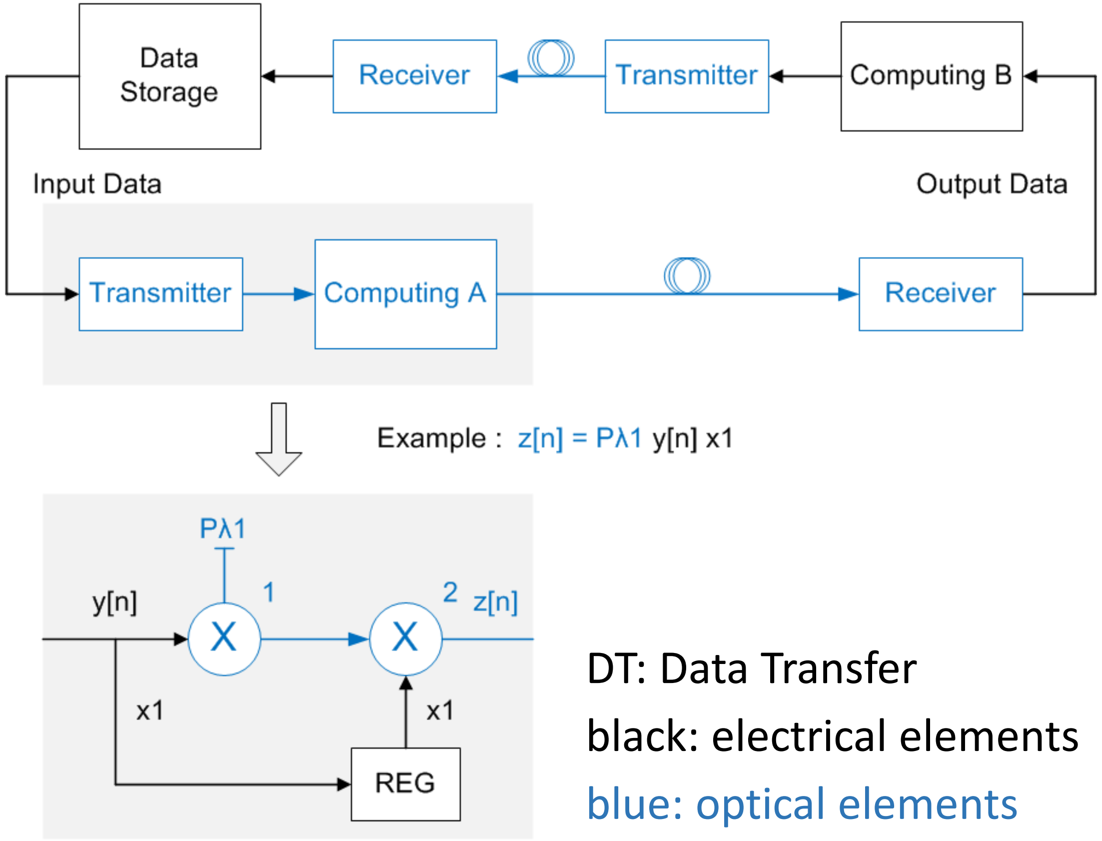

DT: Data Transfer

black: electrical elements

blue: optical elements

Data Transfer Compute Model Optimized for Math Operations

Electrical Computing w/ optical DT

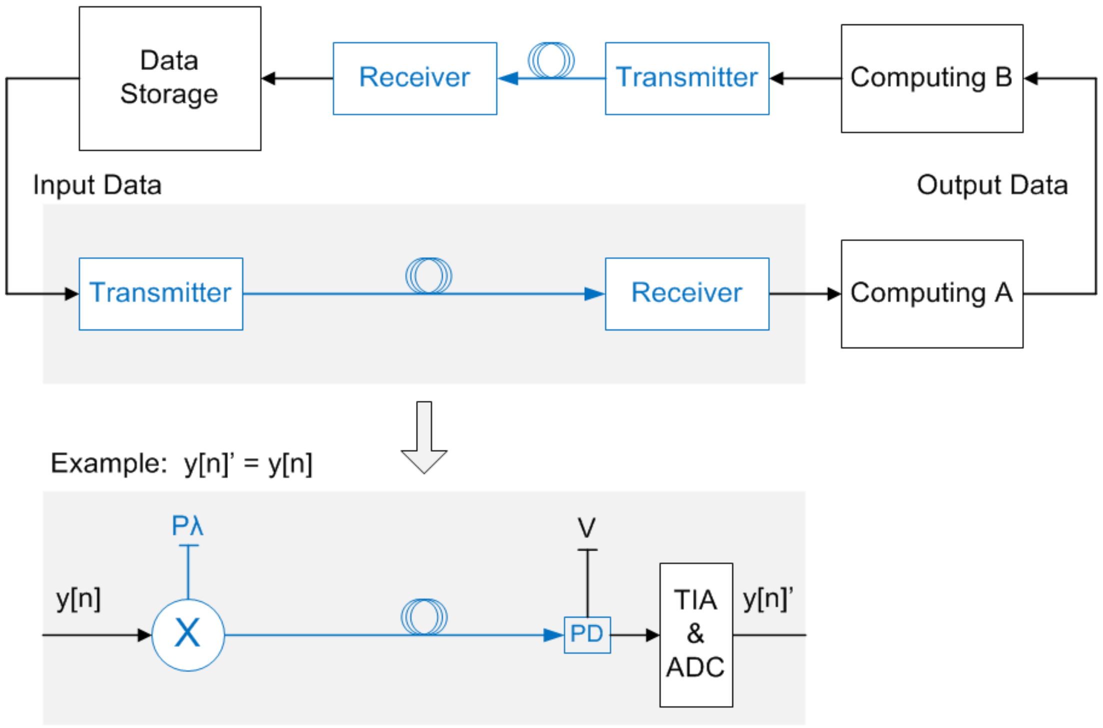

DT: Data Transfer

black: electrical elements

blue: optical elements

Apples-to-Apples Energy Use Comparison Models

Electrical Computing w/ optical DT

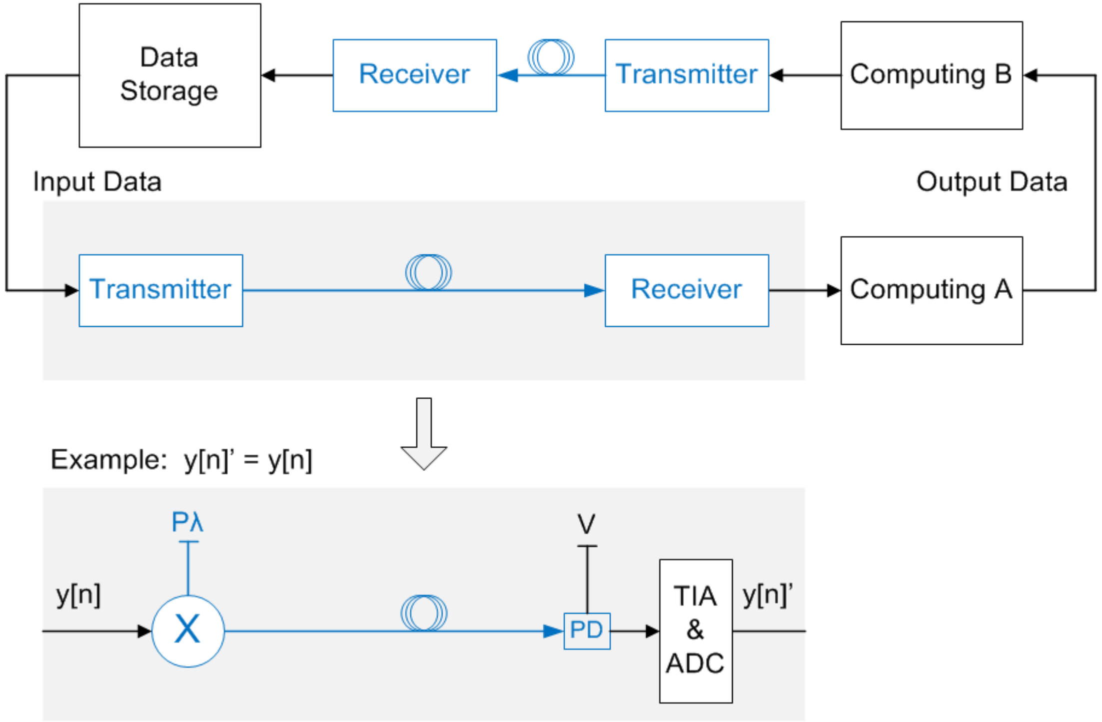

Data Storage

Receiver

Transmitter

Computing B

Input Data

Output Data

Transmitter

Receiver

Computing A

Example: $y[n]' = y[n]$

PλV

y[n]

X

PD

TIA & ADC

y[n]'

Optical Computing w/ optical DT

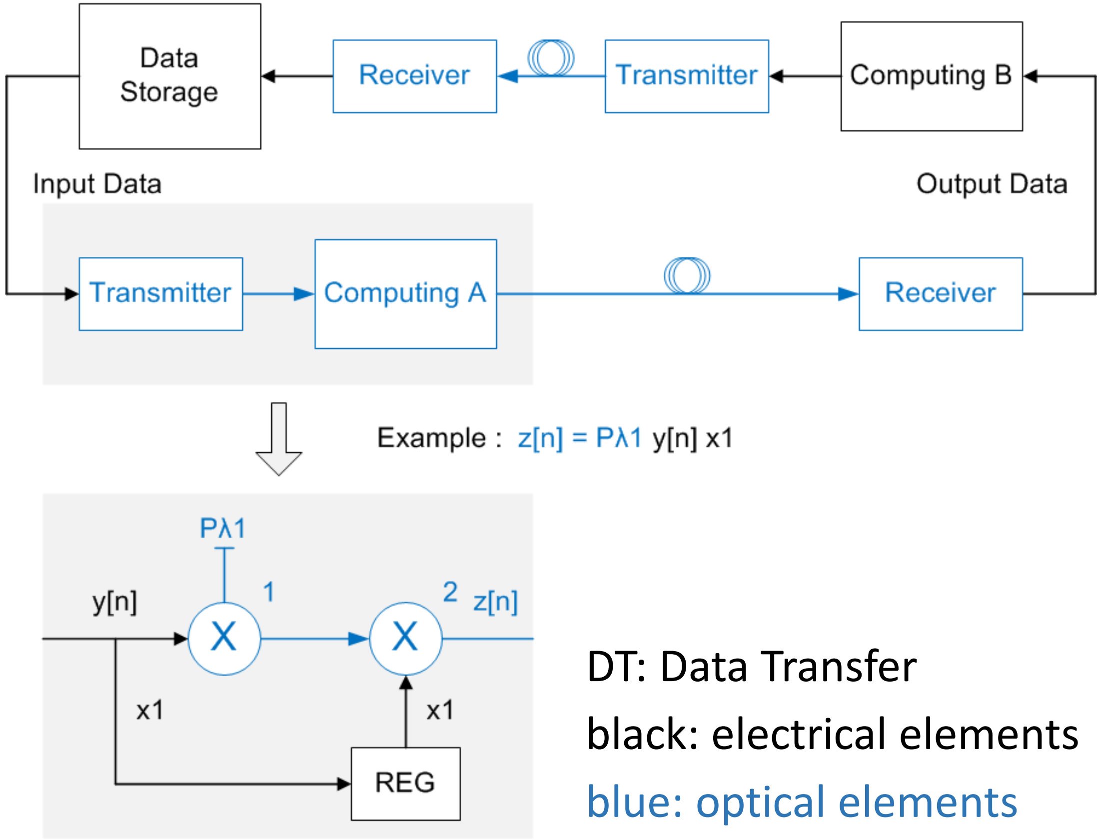

Data Storage

Receiver

Transmitter

Computing B

Input Data

Output Data

Transmitter

Computing A

Receiver

Example : $z[n] = P\lambda 1\ y[n]\ x1$

Pλ1

y[n]

1

X

2

z[n]

X

x1

x1

REG

DT: Data Transfer

black: electrical elements

blue: optical elements

Outline

- Optics & AI
- Optics Modulation Alternatives
- Optical Computing
  - Frenzy Examples
  - Historical Examples
  - Precision
  - Comparison
  - ➢**Addition**
  - Multiplication
  - Convolution
  - Discussion
- Optics: Reality or Fantasy

Electrical Parallel Addition (column sum) Model

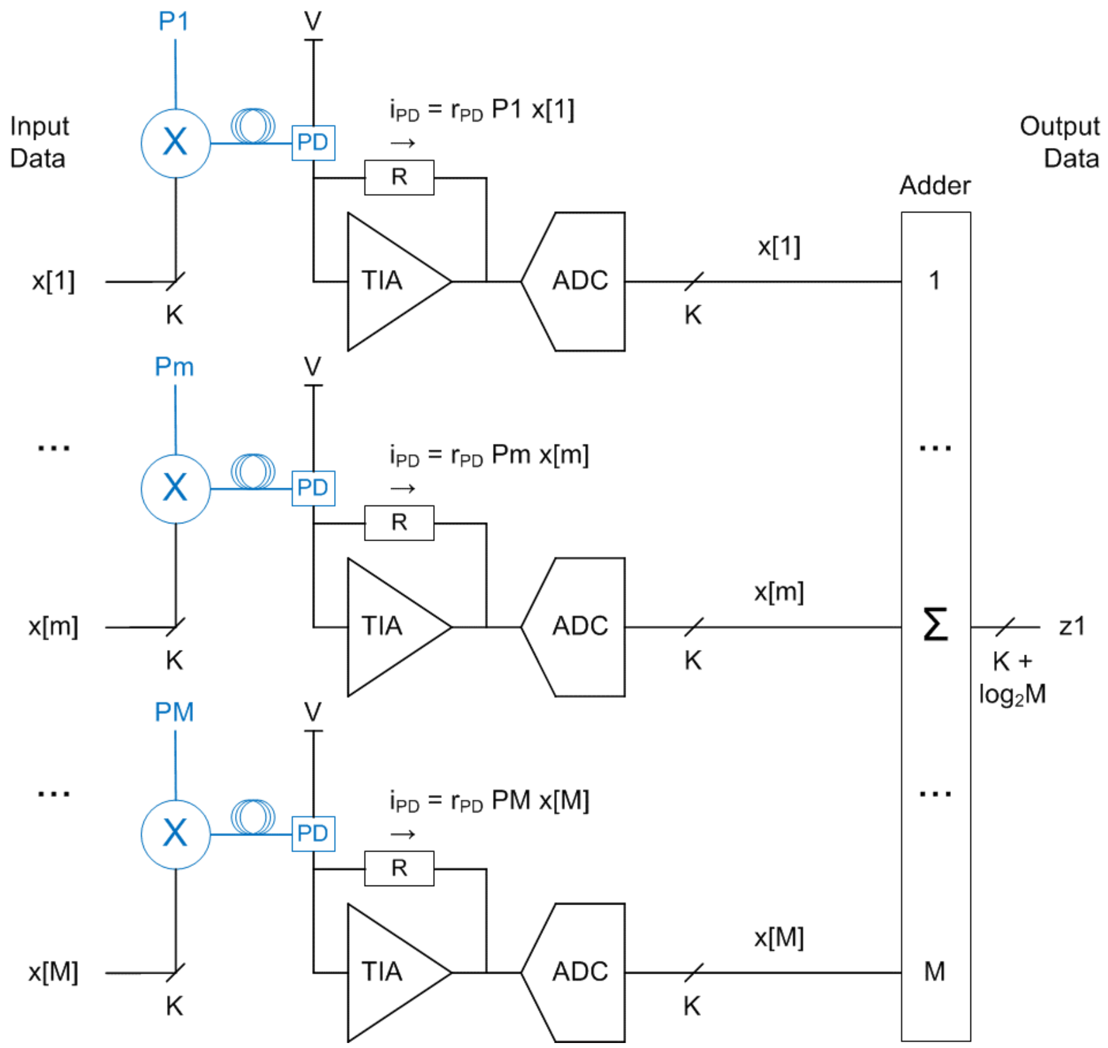

black: electrical elements

blue: optical elements

Electrical and Optical Parallel Addition (column sum) Models

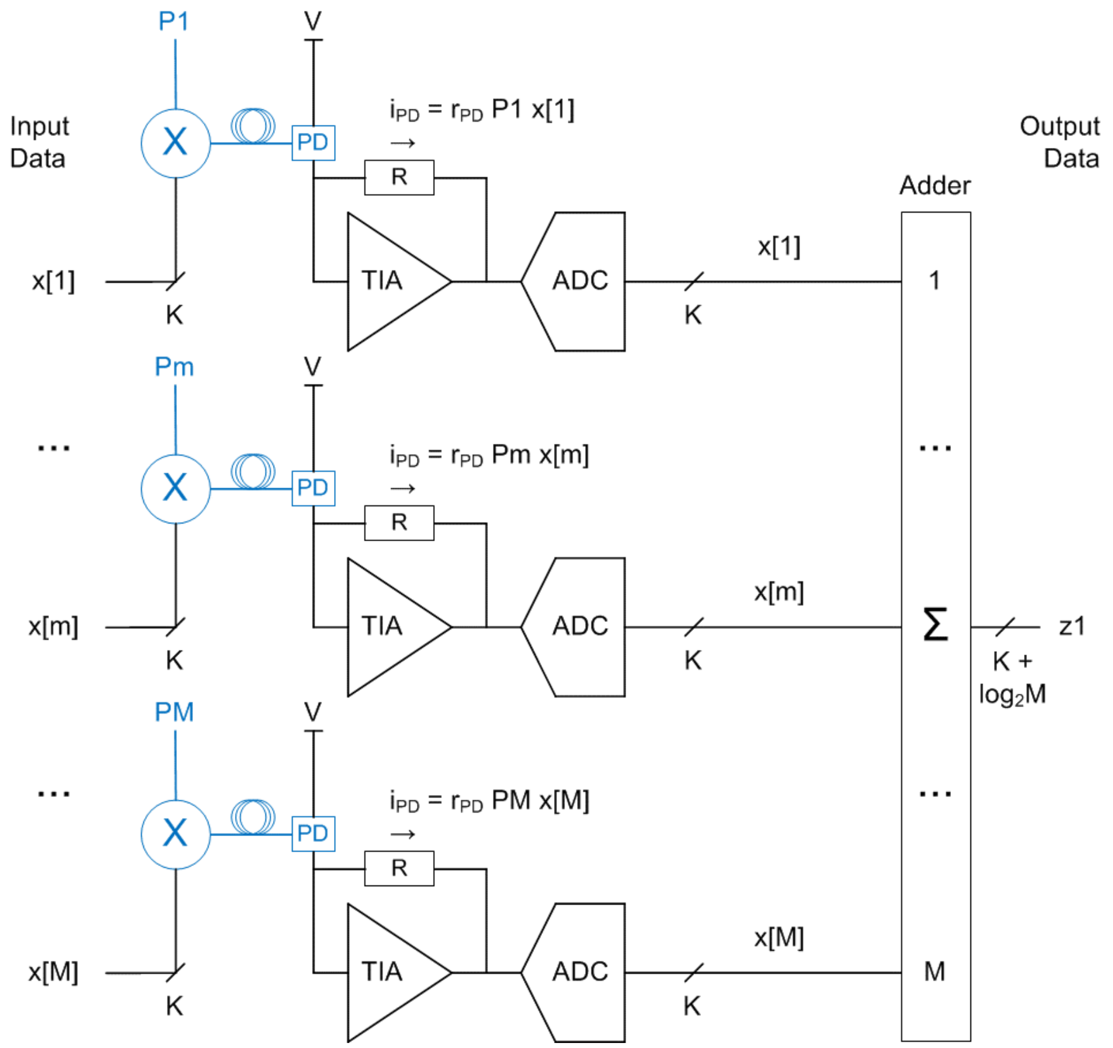

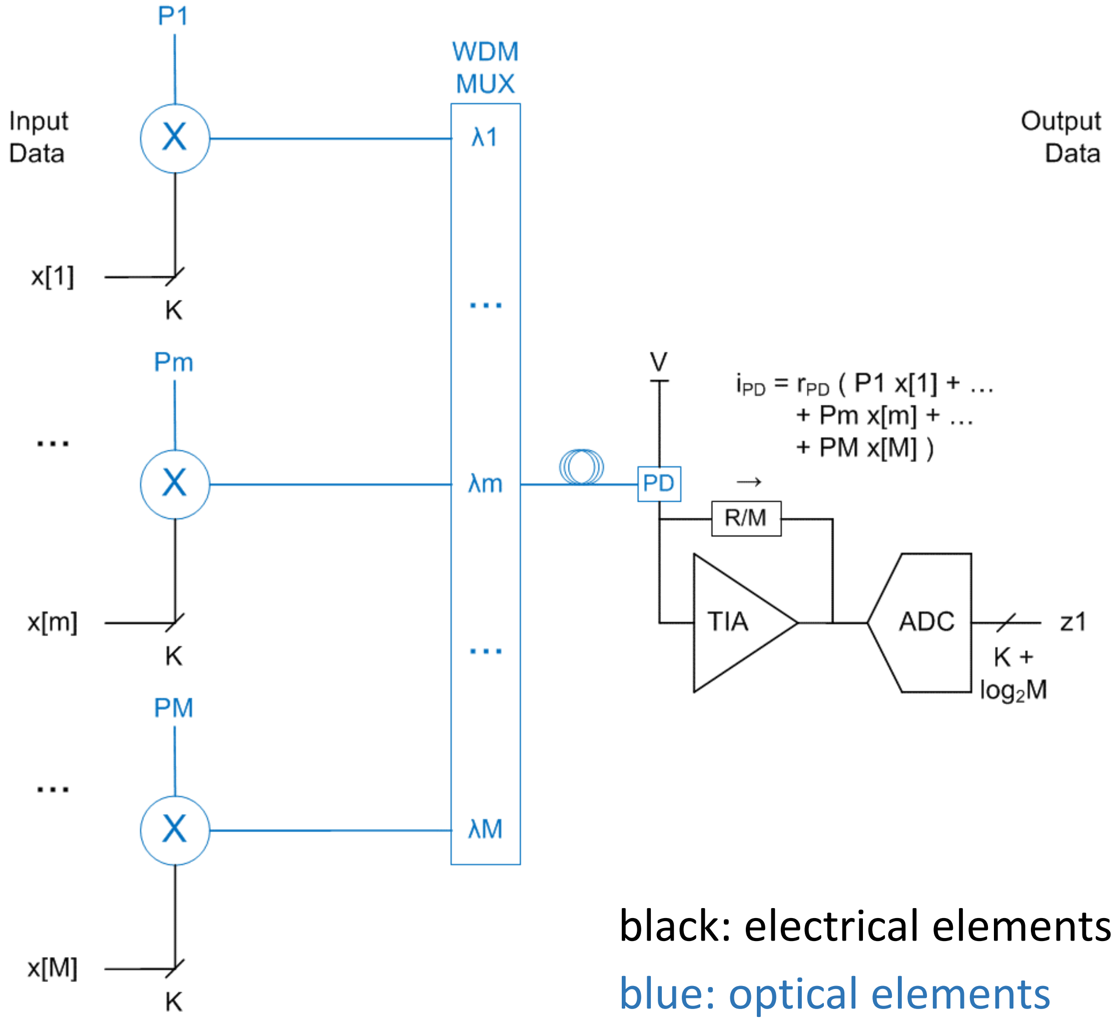

black: electrical elements

blue: optical elements

CMOS Adder Energy Use

CMOS node	Delay	Energy/op (max)	Input	Rate	Energy
nm	ps	fJ	bits/op	Gops/s	fJ/bit
7	40	50	16	25	2.9
7	30	40	16	33	2.5
				average	**2.7**

Q. Xie, X. Lin, S. Chen, M. Dousti and M. Pedram, "Performance Comparisons between 7nm FinFET and Conventional Bulk CMOS Standard Cell Libraries," *IEEE Transactions on Circuits and Systems II: Express Briefs,* vol. 62, no. 8, pp. 761-765, August 2015.

A. Vatanjou, E. Lte, T. Ytterdal and S. Aunet, "Ultra-low Voltage and Energy Efficient Adders in 28nm FDSOI Exploring Poly-biasing for Device Sizing," Microprocessors & Microsystems, vol. 56, no. C, pp. 92-100, February 2018.

A. Stillmaker and B. Baas, "Scaling equations for the accurate prediction of CMOS device performance," *Integration the VLSI journal,* vol. 58, pp. 74-81, February 2017.

Energy Use of High-speed CMOS ADCs

Output	**Rate**	**CMOS node**	**Effective bits**	**Energy**	**Reference**
Bits	GS/s	nm	ENOB	**fJ/bit**	
6	24	28	4.5	**210**	[32]
6	3.3	28	5.4	310	[33]
8	10	65	6.4	800	[34]
8	1	28	7.3	350	[35]
8	28	7	5.0	355	[36]

References from C. Cole, "Optical and electrical programmable computing energy use comparison," Optics Express, Vol. 29, Issue 9, pp. 13153-13170, 2021.

Electrical & Optical Addition Energy Use Comparison

- 90 Gops/sec 8-bit 7nm CMOS adder energy use:
  - **<2%** that of 28 Gops/sec 8-bit 7nm CMOS ADC
- M K-bit ADCs (electrical computation) energy use
  - **equals** one (K + $\log_2 M$)-bit ADC (optical computation)
- ADC use dominates CMOS and Optical addition energy use

Outline

- Optics & AI
- Optics Modulation Alternatives
- Optical Computing
  - Frenzy Examples
  - Historical Examples
  - Precision
  - Comparison
  - Addition
  - ➢**Multiplication**
  - Convolution
  - Discussion
- Optics: Reality or Fantasy

Electrical Parallel Vector Multiplication Model

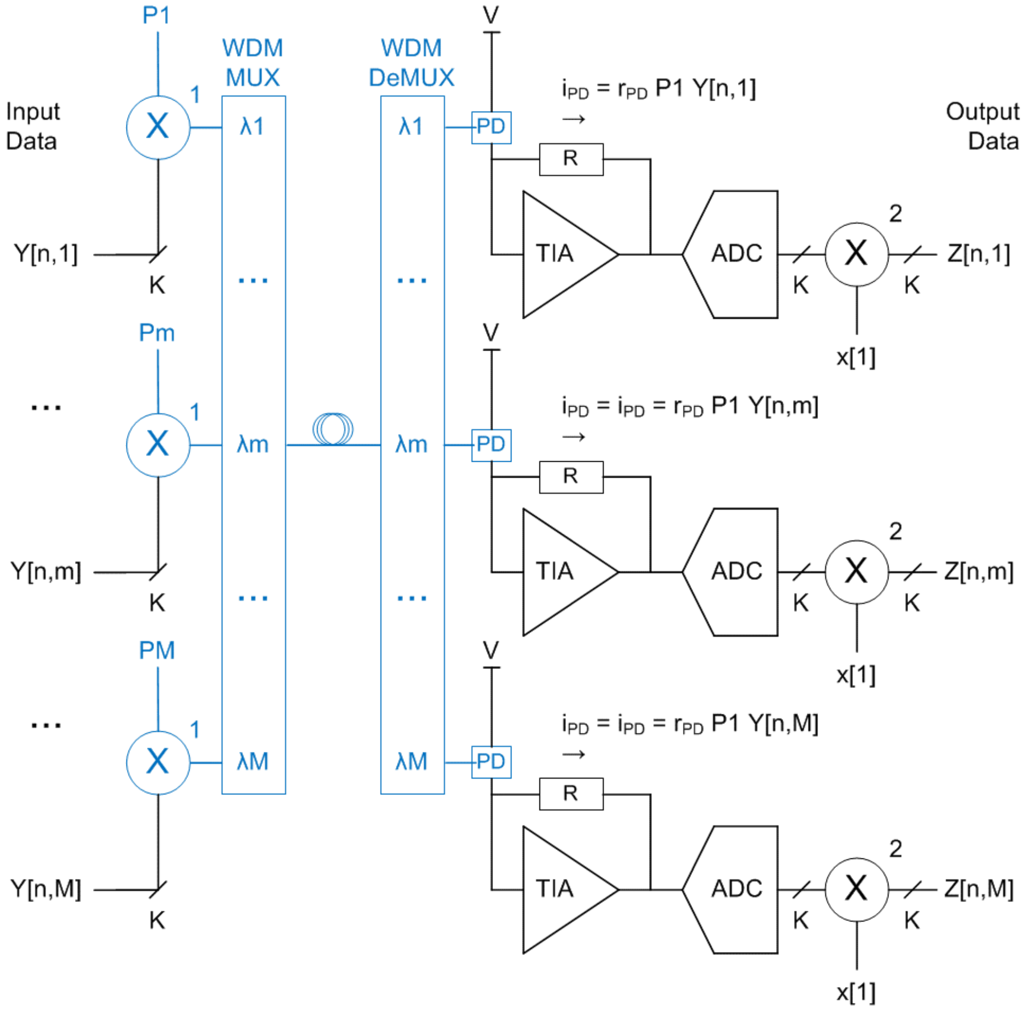

Electrical and Optical Parallel Vector Multiplication Models

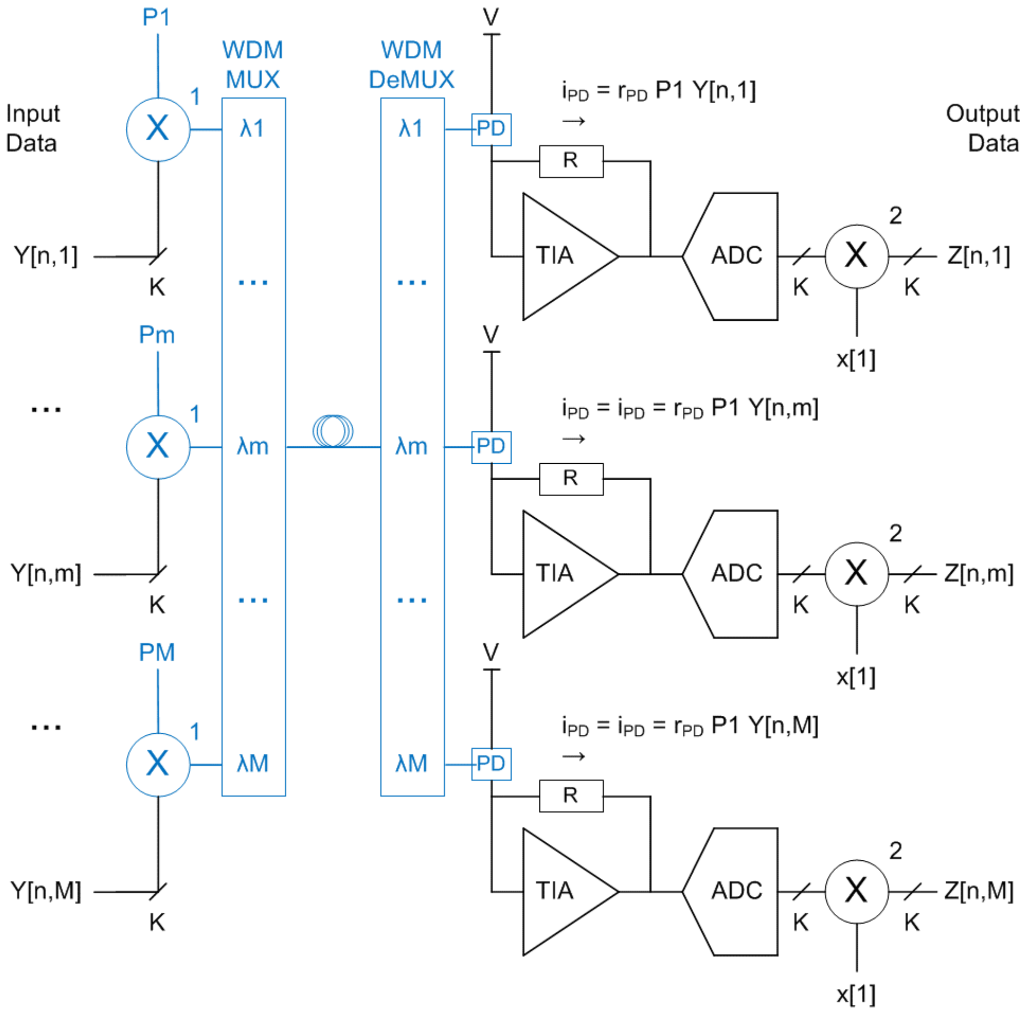

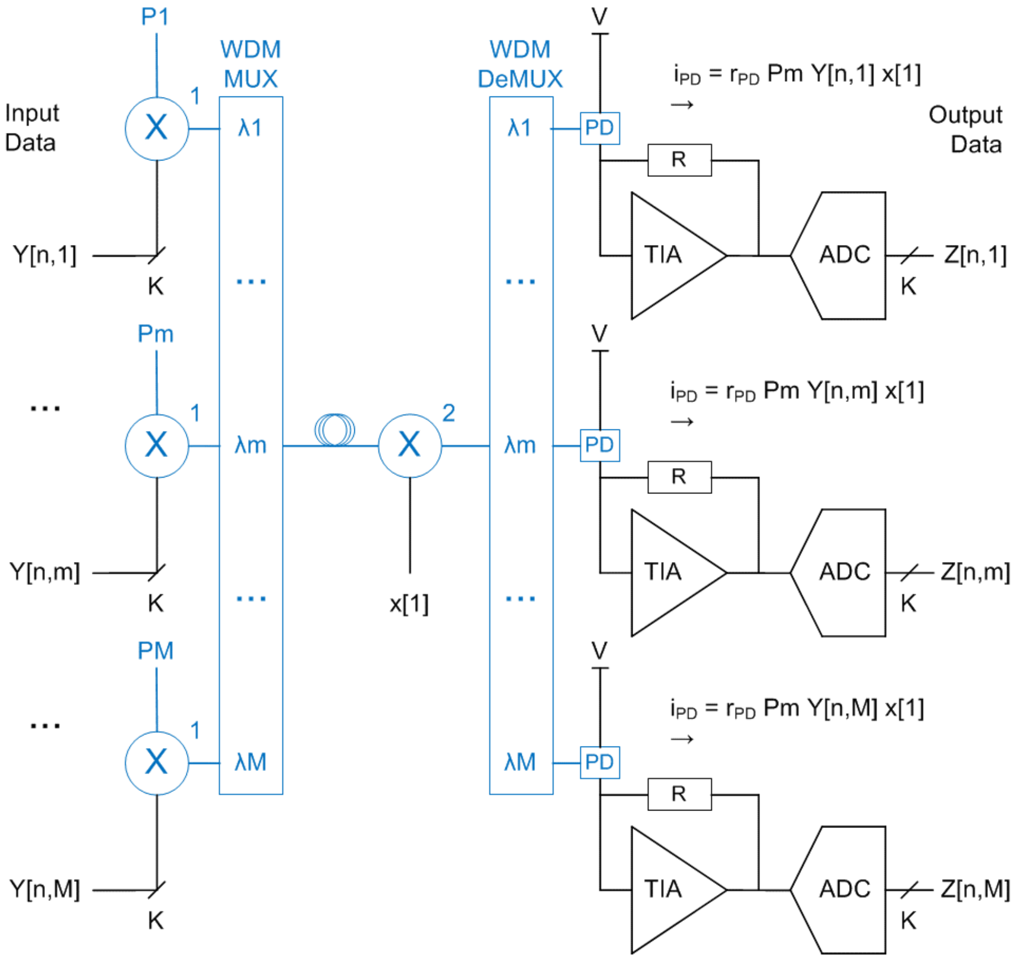

Energy Use of CMOS 16-bit Multipliers

CMOS node	Delay	Energy/op (max)	Input	Rate	Energy
nm	ps	fJ	bits/op	Gops/s	fJ/bit
7	58	296	16	17.5	19
7	40	310	16	25	19

Q. Xie, X. Lin, S. Chen, M. Dousti and M. Pedram, "Performance Comparisons between 7nm FinFET and Conventional Bulk CMOS Standard Cell Libraries," *IEEE Transactions on Circuits and Systems II: Express Briefs,* vol. 62, no. 8, pp. 761-765, August 2015.

D. Baran, M. Aktan, and V. Oklobdzija, "Energy Efficient Implementation of Parallel CMOS Multipliers with Improved Compressors," in ACM/IEEE International Symposium on Low-Power Electronics and Design (ISLPED), pp. 147–152, August 2010.

A. Stillmaker and B. Baas, "Scaling equations for the accurate prediction of CMOS device performance," *Integration the VLSI journal,* vol. 58, pp. 74-81, February 2017.

Electrical & Optical Multiplication Energy Use Comparison

- 25 Gops/sec 16-bit 7nm CMOS multiplier energy use:
  - **9%** that of 28 Gops/sec 8-bit 7nm CMOS ADC
- ADC use dominates CMOS and Optical multiplication energy use

Outline

- Optics & AI
- Optics Modulation Alternatives
- Optical Computing
  - Frenzy Examples
  - Historical Examples
  - Precision
  - Comparison
  - Addition
  - Multiplication
  - ➢**Convolution**
  - Discussion
- Optics: Reality or Fantasy

Electrical Matrix Vector Product Computation Model

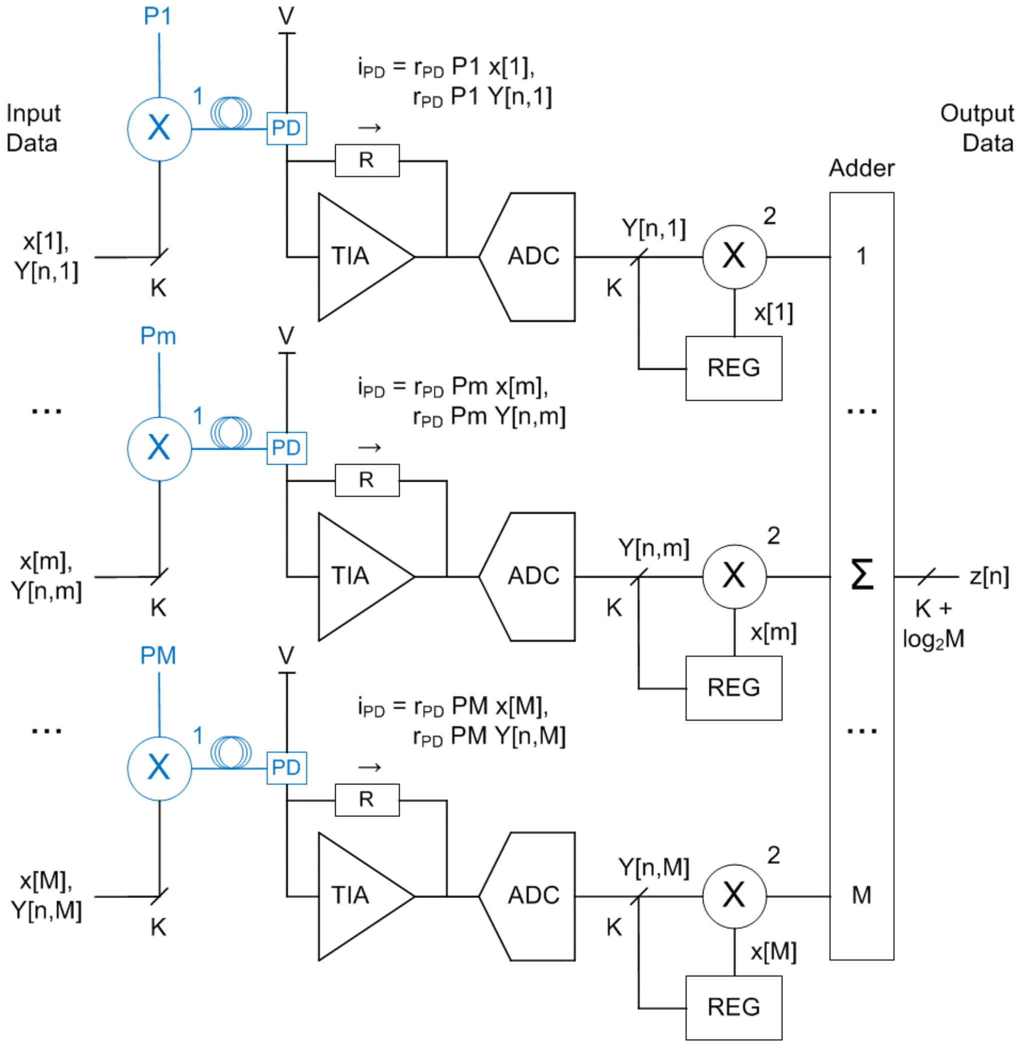

$i_{PD} = r_{PD}$ P1 x[1], r_{PD} P1 Y[n,1]

$i_{PD} = r_{PD}$ Pm x[m], r_{PD} Pm Y[n,m]

$i_{PD} = r_{PD}$ PM x[M], r_{PD} PM Y[n,M]

Electrical & Optical Matrix Vector Product Computation Models

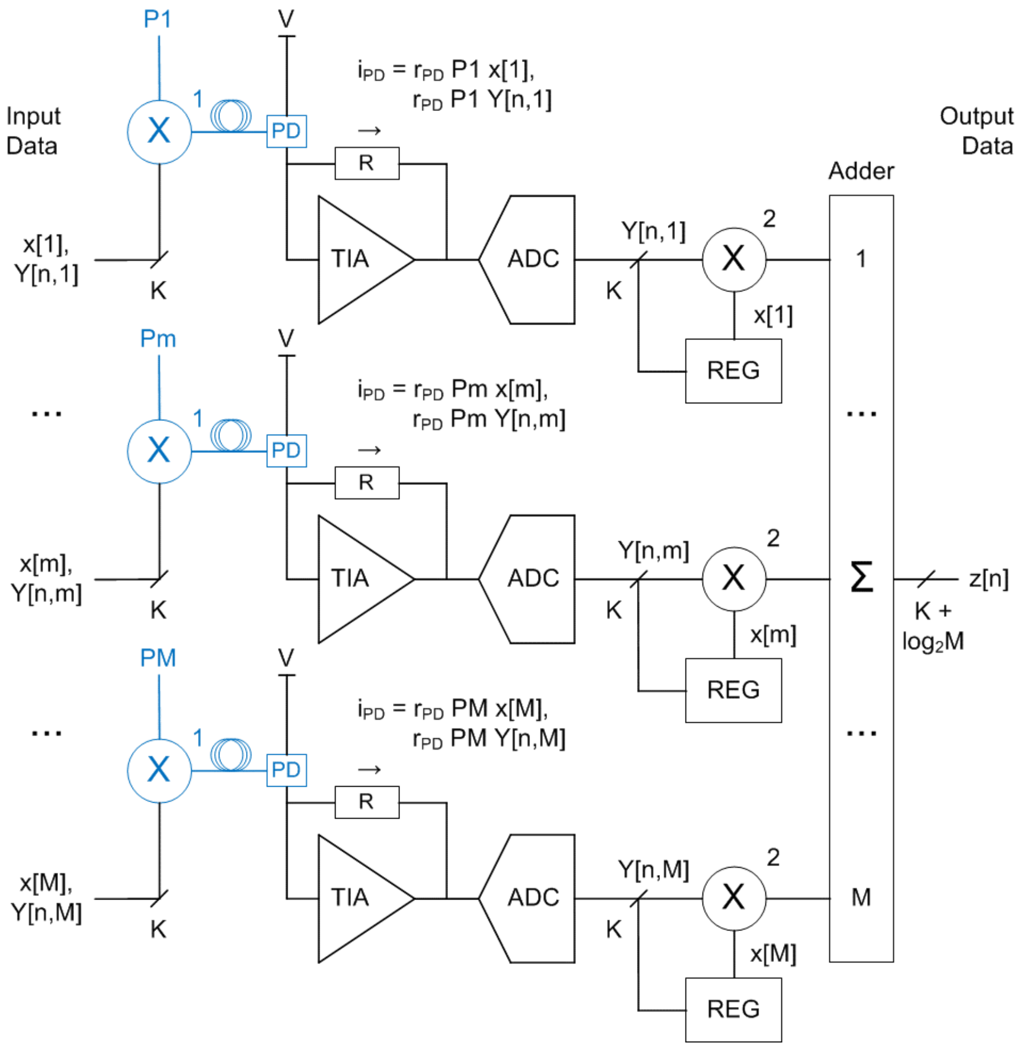

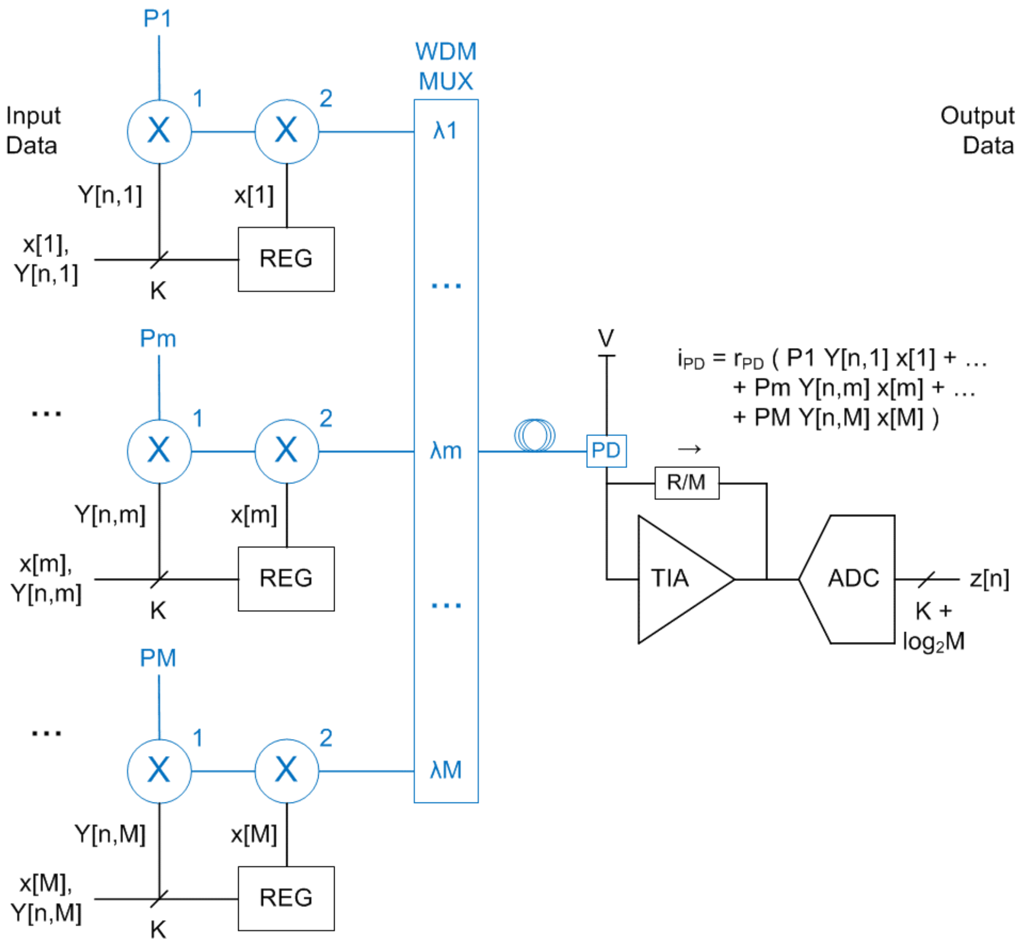

CMOS Multiplier and FFE MAC Energy Use

MAC Type	CMOS node	Delay	Energy/op (max)	Input	Rate	Energy
	nm	ps	fJ	bits/op	Gops/s	fJ/bit
Adder & Multiplier	7	58	367	16	17.5	23
FFE	7	11	159	8	90	**20**

Q. Xie, X. Lin, S. Chen, M. Dousti and M. Pedram, "Performance Comparisons between 7nm FinFET and Conventional Bulk CMOS Standard Cell Libraries," *IEEE Transactions on Circuits and Systems II: Express Briefs,* vol. 62, no. 8, pp. 761-765, August 2015.

C. Menolfi, M. Braendli, P. Francese, T. Morf, A. Cevrero, M. Kossel, L. Kull, D. Luu, I. Ozkaya and T. Toifl, "A 112Gb/s 2.6pJ/b 8-tap FFE PAM-4 SST TX in 14nm CMOS," in IEEE International Solid-State Circuits Conference Digest of Technical Papers, pp. 104-105, February 2018.

A. Stillmaker and B. Baas, "Scaling equations for the accurate prediction of CMOS device performance," *Integration the VLSI journal,* vol. 58, pp. 74-81, February 2017.

Electrical & Optical Convolution Energy Use Comparison

- 90 Gops/sec 8-bit 7nm CMOS convolution MAC energy use:
  - **10%** of 28 Gops/sec 8-bit 7nm CMOS ADC energy use
- M K-bit ADCs (electrical computation) energy use
  - **equals** one (K + $\log_2 M$)-bit ADC (optical computation)
- ADC use dominates CMOS and Optical convolution energy use

Optical Datacom Filter CMOS Computing Example: Fast FFE

- FFE (Feed Forward Equalizer) processing is convolution
- Same processing as applying weights in a neuron (inference)
- Used in high volume PHY (CDR) optical receivers
- Architecture: ADC + CMOS DSP with only CTLE analog pre-compensation
- Optical receiver FFE is the perfect problem for optical computing:
  - high bit rate
  - low precision
  - low number of coefficients
  - digital to optical & optical to digital conversion already in place
  - zero incremental energy use
- Yet all optical receivers use CMOS DSP FFEs, and none use optical computing

Outline

- Optics & AI
- Optics Modulation Alternatives
- Optical Computing (ref: https://opg.optica.org/oe/fulltext.cfm?uri=oe-29-9-13153)
  - Frenzy Examples
  - Historical Examples
  - Precision
  - Comparison
  - Addition
  - Multiplication
  - Convolution
  - ➢**Discussion**
- Optics: Reality or Fantasy

General Purpose Optical Computing

- Low-precision; INT4 (despite claims to the contrary)
  - not useful standalone in the Datacenter
- No energy use savings over digital CMOS HW
  - dominated by Processor ⟺ Memory data movement
- Size ~ 1K x greater per dimension:
  - 1M x greater chip area
- Optical Compute HW is not and will not be used in Datacenters

Pre-processing Optical Computing

- Examples:
  - Eyeglasses
  - Digital camera front-end
  - LIDAR
  - FSO Beamforming
  - Microscopy
- Great R&D and Product area

There is Hope for General Purpose Optical Computing

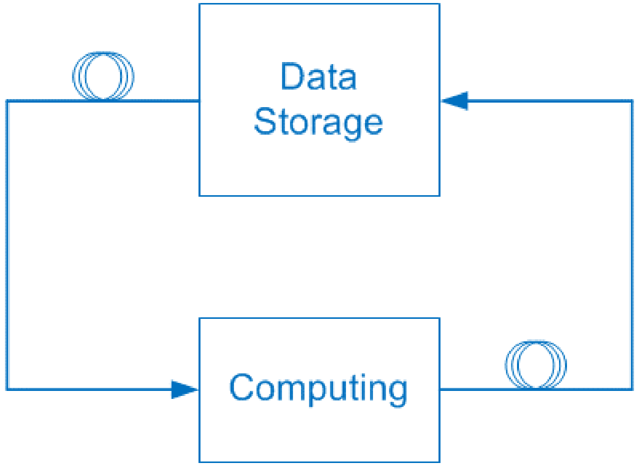

blue: optical elements

It just needs dense optical RAM.

Outline

- Optics & AI
- Optics Modulation Alternatives
- Optical Computing
  - Frenzy Examples
  - Historical Examples
  - Precision
  - Comparison
  - Addition
  - Multiplication
  - Convolution
  - Discussion
- **Optics: Reality or Fantasy**

Optics: Reality

Technical data

- Shannon Hartley: $C = B \log_2(1 + S/N)$
  - measured full link BER curves →
- Reliability: System MTBF or FIT
  - accelerated aging data calculations
- Multi-channel: WDM, TDM, Spatial, other
  - BER with all channels ON
  - reliability of all channels
- Integrated: CPO, Chiplets, Wafer-scale
  - BER with entire system ON, over all process and temp corners
  - reliability data of the entire system

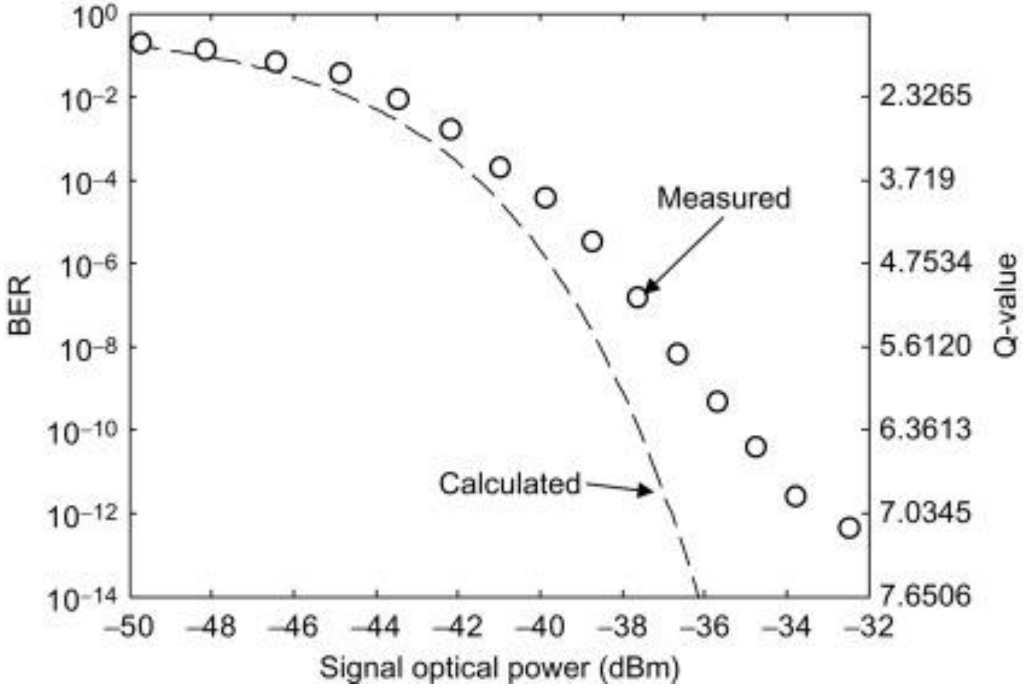

Optics: Fantasy

Marketing claims

- In AI normal engineering doesn’t apply
- 10x to 100x performance improvement
- Application only verification
- No statistically valid measured link BER curves
- No system FIT data calculations
- Individual component data best case

Reality and Fantasy in Optics

Thank you