



ORACLE[®]

Optical interconnect circuits: some design considerations

Frankie Liu, Ph.D.
VLSI Research Group
Sun Labs, Oracle

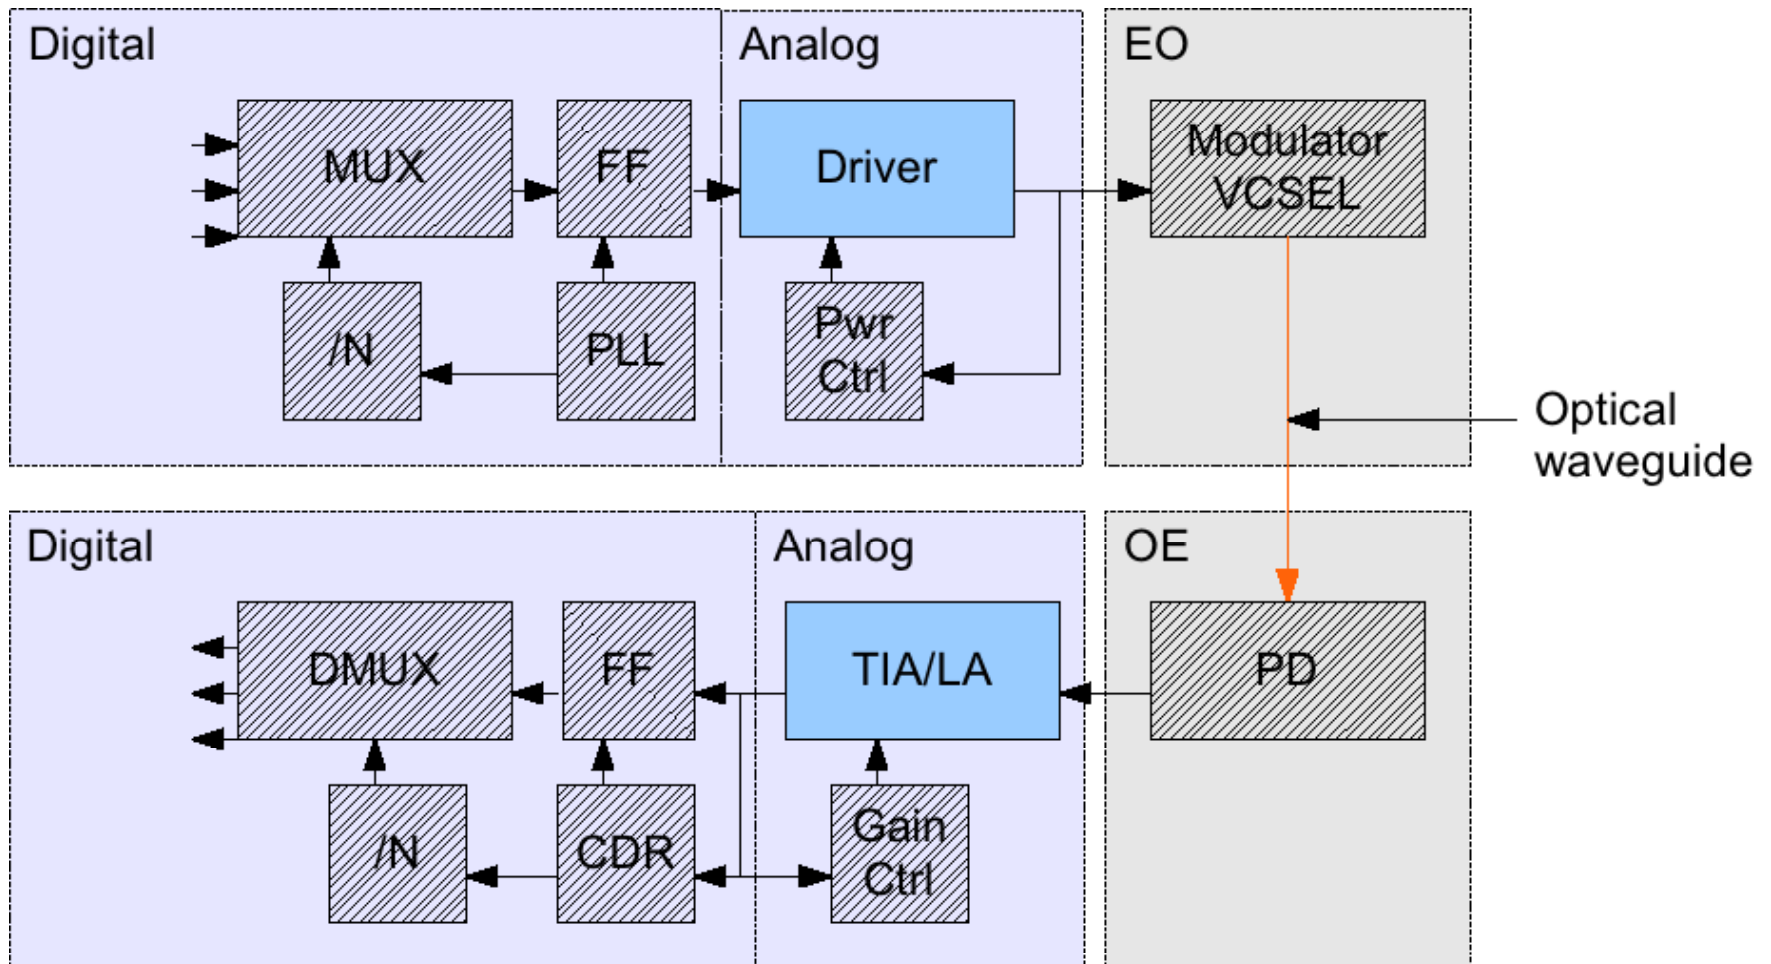


Why consider optics?

- Increased interest in silicon photonic interconnect
 - Potential win in density, energy, BW over wires
- From an architectural view, optics is just another tool
 - Use it if its characteristics give a net “win”
 - But then we need to understand its costs
- This is an introduction to these costs
 - From a circuit/physical design perspective
 - A list of things to consider in a system evaluation

A generic serialized system

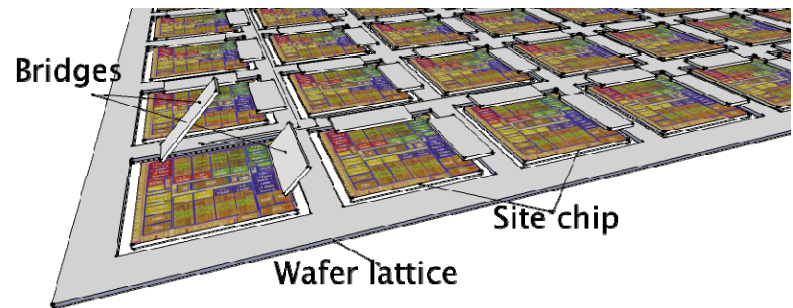
Here we consider just part of the whole system



Designers tweak all parts of this system

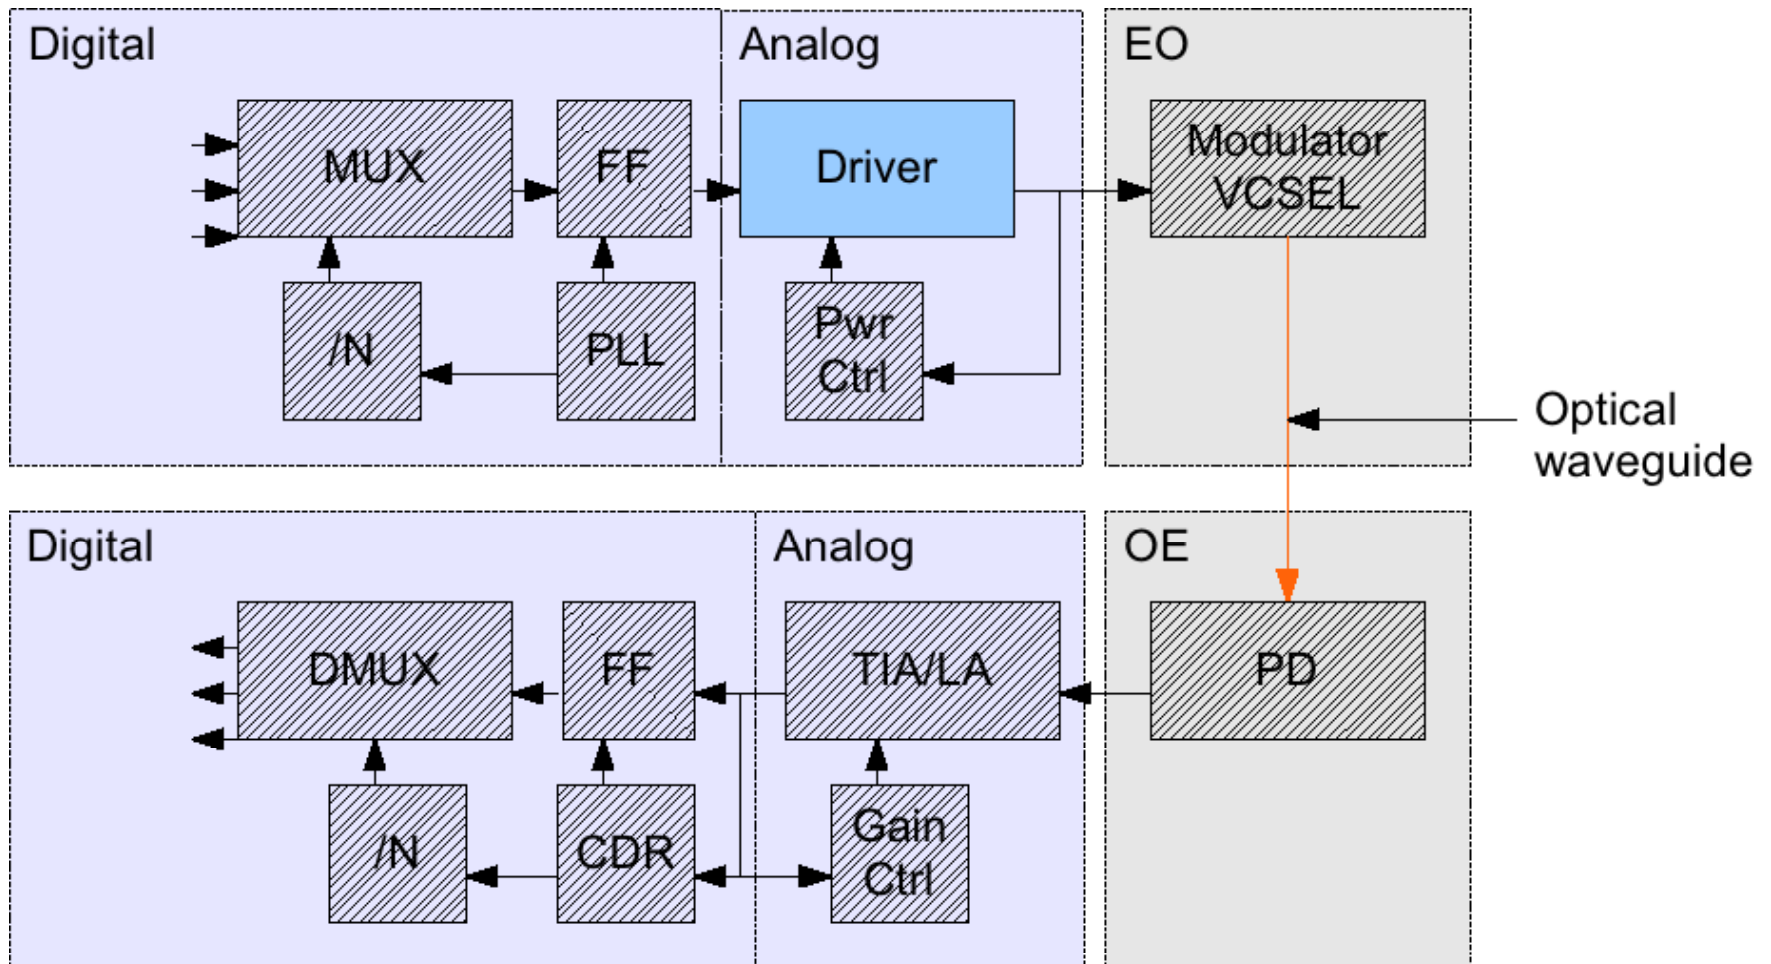
Lots of design effort in the CDR, PLL, mux/demux, etc.

- Link design depends on the architectural use-case
 - Some applications allow system simplifications
- E.g., Oracle's "macrochip"
 - Co-packaged chips sharing a silicon substrate [8]
 - Optical links are inter-chip, but within-package
 - Shared mesochronous clock between RX/TX
 - Use periodic calibration (common time-sense)
 - DWDM has potential for wide parallel I/O




But today we'll just cover a few simple themes

Start with the driver



Transmitter basics

- TX converts electrical signals into optical signals
- Many device flavors
 - Local light sources
 - VCSELs, LEDs
 - CMOS integration tricky
 - Modulating a remote source
 - MZs, rings, quantum wells
 - Better integration prospects



Transmitter basics

- TX converts electrical signals into optical signals

- Many device flavors

- Local light sources

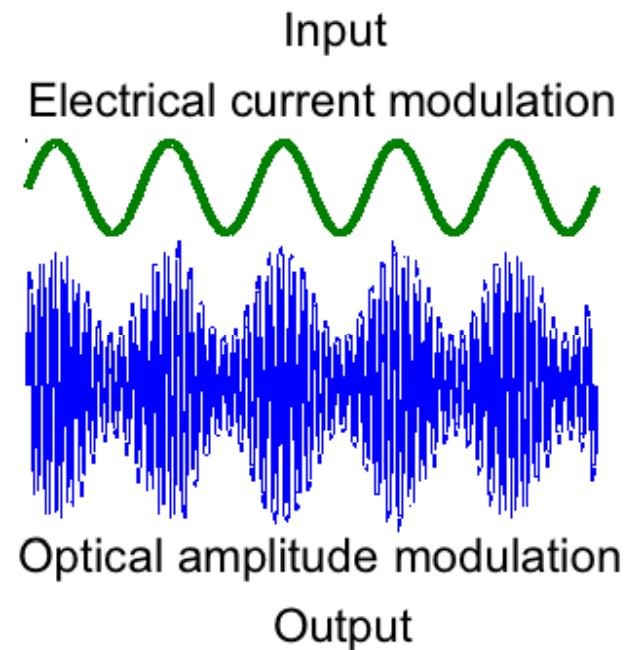
- VCSELs, LEDs

- CMOS integration tricky

- Modulating a remote source

- MZs, rings, quantum wells

- Better integration prospects

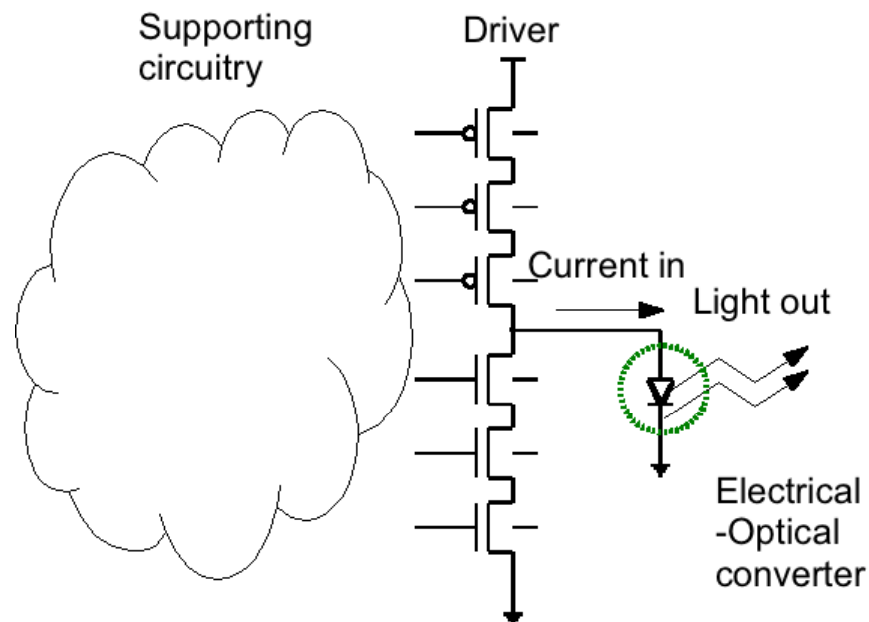


Consider these here

Transmitter basics, con't

Driving ring modulators

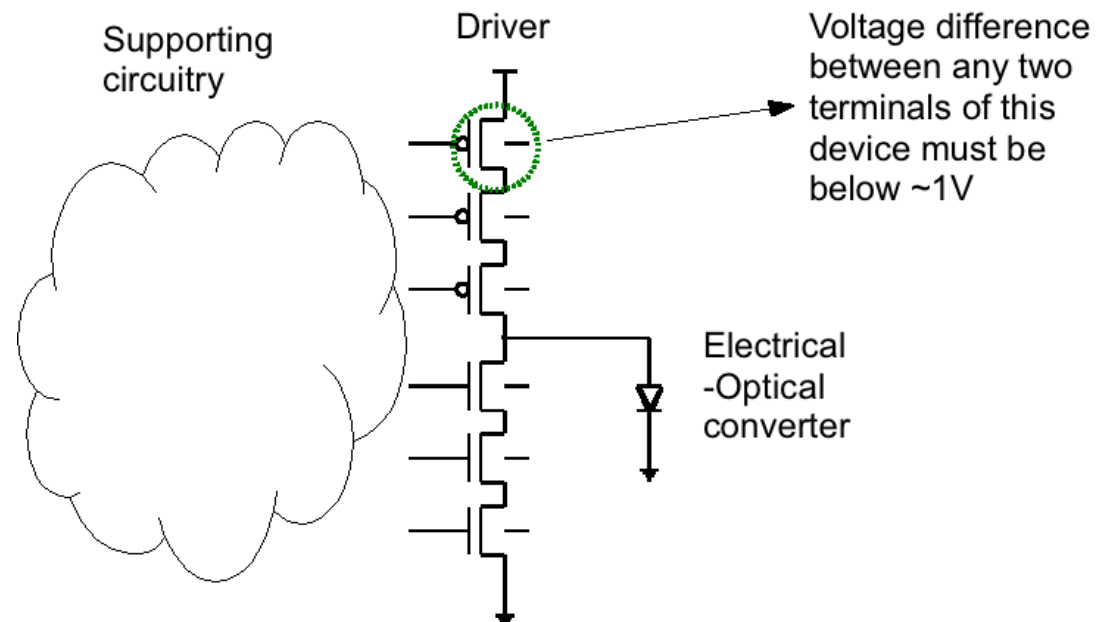
- Good: ring modulators look like a lumped capacitance
 - Simple abstraction for designers; matches CMOS
- Bad: they typically require a high voltage [1]



Transmitter basics, con't

Driving ring modulators

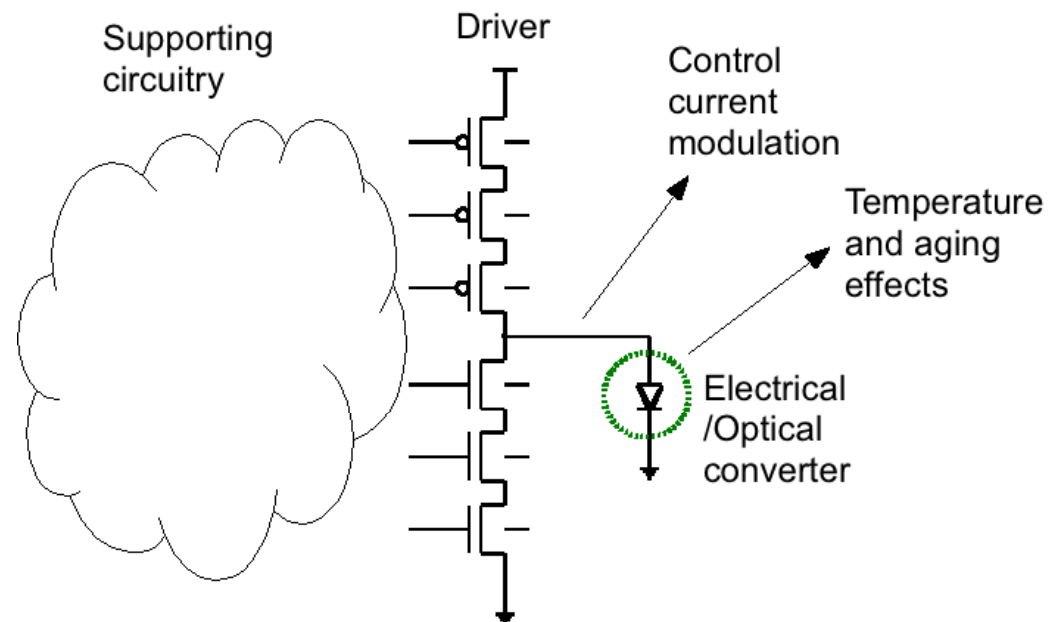
- Fast high voltage switching using low-voltage devices
 - Must take care to protect circuits [2]



Transmitter basics, con't

Driving ring modulators

- Bad: modulators see temperature and aging effects
 - Need driver power and/or environment control [3]






Transmitters and ring modulators

- High-voltage drive is not a very difficult problem
 - At least, it's fairly well understood
- The hard problem is temperature control
 - Rings sensitive to sub-1° K temperature changes
 - Standard solution is to dynamically heat them
 - Is this efficient enough to make rings useful?
 - Until shown (at scale), this is *the* gating issue

Back to the generic system diagram

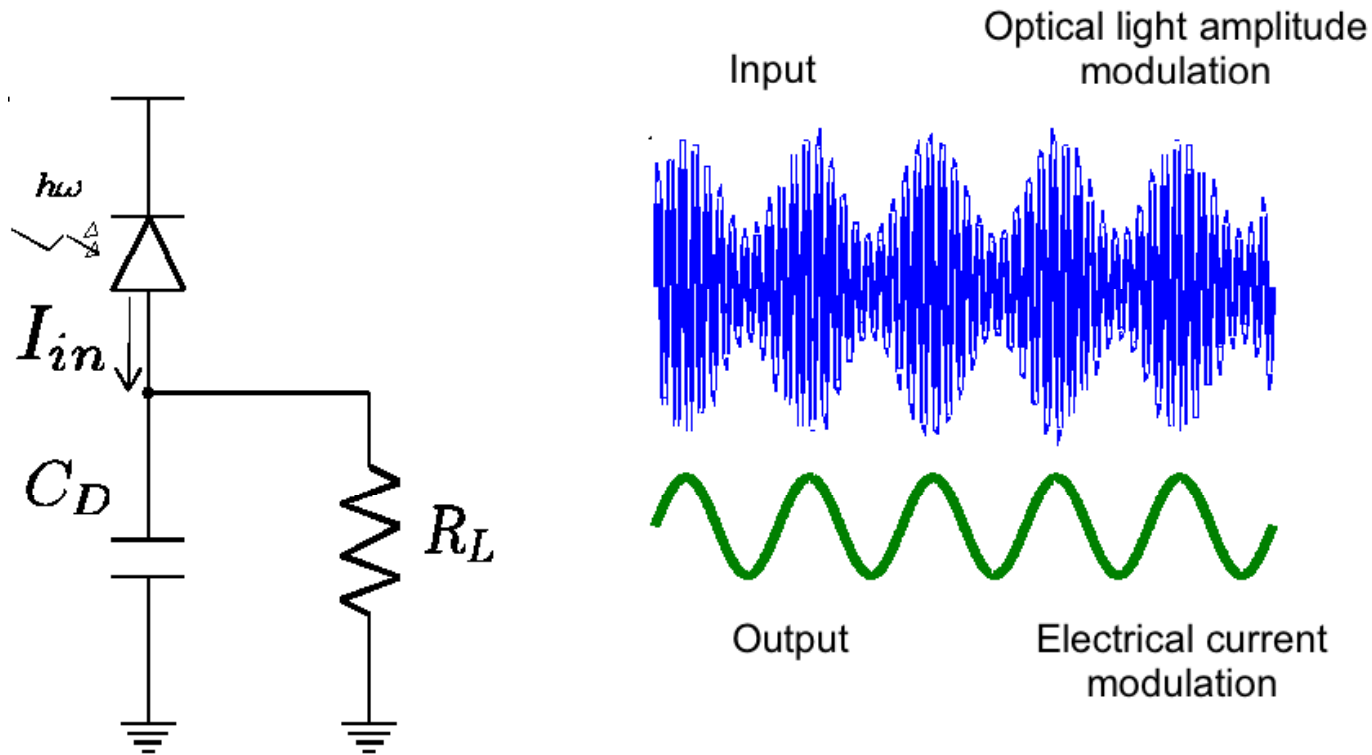
The RX has a different (perhaps more interesting?) set of issues



Receiver basics

Let's start with the simplest system possible


- Simple mechanism: photodiode turns light \rightarrow current
 - “Responsivity” is on order of 1 μA per μW



Not surprisingly, this is a simple circuit

Go back to Circuits 101

- Resistor performs a current-to-voltage conversion
 - Capacitor introduces a low-pass filter pole [5]



A key characteristic of any receiver is its SNR

Signal-to-noise ratio

- For this receiver, SNR is relatively simple

$$BW = \frac{1}{R_L C_D}$$


$$\text{Gain} = R_L$$

$$\text{Noise} = \frac{kT}{C_D}$$

$$\text{SNR} = \frac{I_{in}^2 R_L^2}{\frac{kT}{C_D}} = \frac{I_{in}^2}{kT C_D BW^2}$$

- Note that SNR is independent of R_L
 - For a fixed BW and C_D


Thermally induced voltage fluctuations



Why do we care about SNR?

(If it were *really* important, wouldn't the internet cease to exist?)

- Larger SNR → smaller probability of error
 - A reasonable estimate of BER: $10^{-\text{SNR}/4.5}$



SNR example

Signal-to-noise ratio

- Let's pick some interesting numbers

$$BW = \frac{1}{R_L C_D}$$


$$\text{Gain} = R_L$$

$$\text{Noise} = \frac{kT}{C_D}$$

$$\text{SNR} = \frac{I_{in}^2 R_L^2}{\frac{kT}{C_D}} = \frac{I_{in}^2}{kT C_D BW^2}$$

$BW = 2\pi \times 4\text{GHz}$	} SNR = 760
$C_D = 50\text{fF}$	
$I_{in} = 10\mu\text{A}$	


- This is huge! So what's the catch?



The catch is that the gain is too small

Cascaded amplifiers (to get a ~1V signal at the output) degrade the SNR

- Noise of the cascade is set by noise of the 1st stage
 - So concentrate on a low noise factor 1st stage





But how do we get more gain in the first place?

Big BW requires small capacitance and small resistance

- Recall that


$$\text{BW} = \frac{1}{R_L C_D} \quad \text{Gain} = R_L$$

- So in this case, R_L (=gain) must be at most 800Ω
 - An input signal of $10\ \mu\text{A}$ turns into at most $8\ \text{mV}$
 - For more signal you need more R_L —and go slower
- R_L tightly couples signal gain and bandwidth
 - How can we decouple these two?
 - Answer: with amplification and feedback

An example first-stage amplifier


A “textbook” transimpedance amplifier (TIA)

- For an equivalent BW, the gain has gone up by A
 - But at what cost? What about noise?



“There ain’t no such thing as a free lunch”

- The amplifier degrades SNR
 - Basic trade-offs independent of amplifier topology



$$\text{SNR} = \frac{I_{in}^2}{kTC_D BW^2} \times f\left(\frac{C_M}{C_D}, \frac{BW}{\omega_T}\right)$$

Relative size of input stage w.r.t. input capacitance

Technology parameter

What does this mean?

$$\text{SNR} = \frac{I_{in}^2}{kTC_D BW^2} \times f\left(\frac{C_M}{C_D}, \frac{BW}{\omega_T}\right)$$

Relative size
of input stage
w.r.t. input
capacitance


Technology
parameter

- SNR relationship shows the fundamental tradeoffs
 - Between BW, input capacitance and input signal
 - This applies to the vast majority of RX topologies
- The function f is of secondary importance
 - Can be derived analytically for simple TIA designs
 - See discussion in [7]

SNR degradation from an inverter amplifier

The actual amplifier topology does not change the results a lot

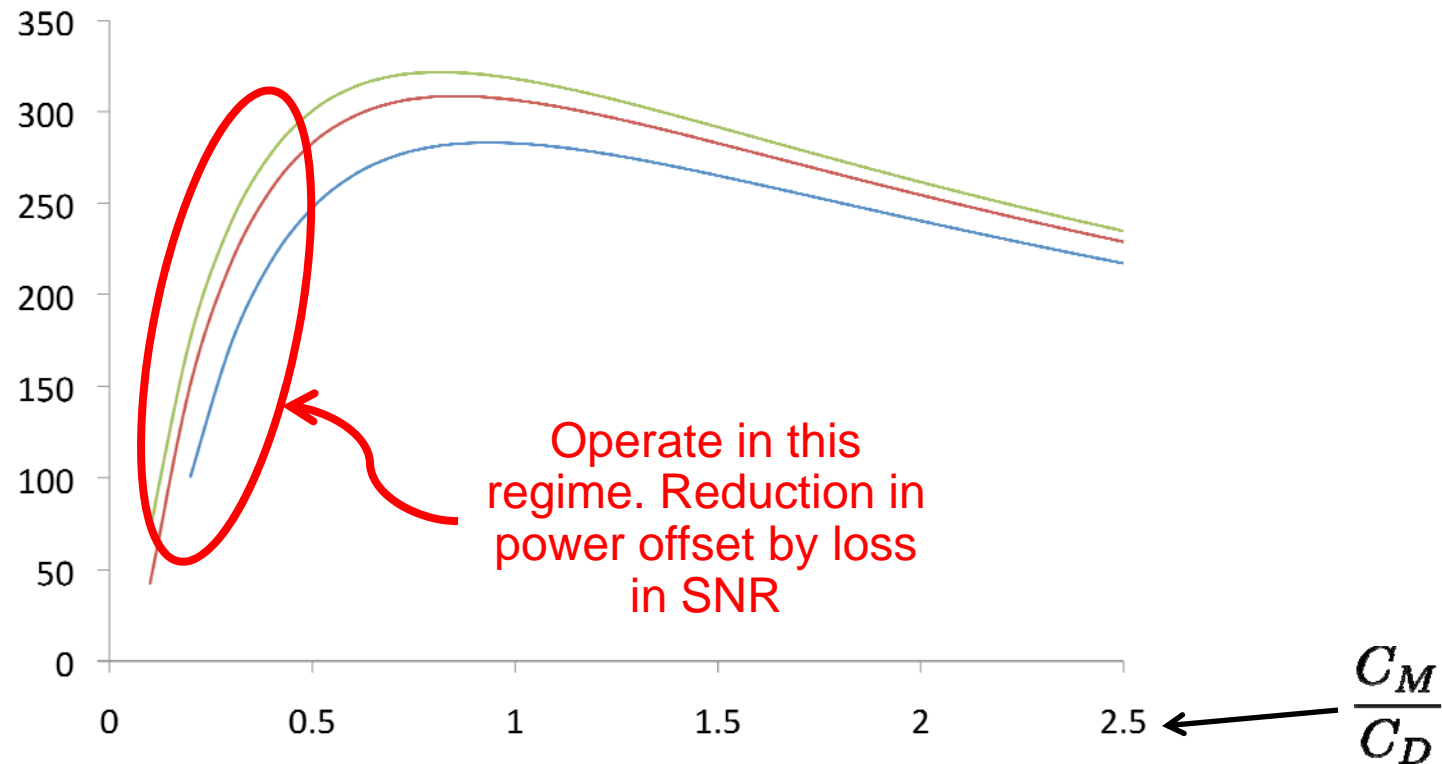
- At optimal $\frac{C_M}{C_D} = \frac{1}{\sqrt{2}}$, SNR degraded by nearly 3X



Operating under power constraints


We don't usually operate at the condition for optimal SNR

- Under power constraints, we pick a required SNR
 - Let that set the power (i.e., the ratio C_M/C_D)



Capacitance ratio C_M/C_D is a proxy for power

- For optimum SNR, the amplifier power is constrained
 - Fixed by the parasitic capacitance C_D



90nm @ 4Gb/s
 SNR=50
 200 fJ/b [6]

$$\text{Power} = \left(\frac{C_M}{C_D} \right)_{opt} \times C_D \times \frac{1\mu\text{m}}{1\text{fF}} \times \frac{150\mu\text{A}}{1\mu\text{m}} \times V_{DD}$$



Bottom line

- TX, RX circuits consume $2/3^{\text{rds}}$ of optical link energy
 - About $1/3^{\text{rd}}$ for each
- TX: Control of temperature and aging effects critical
- RX: For a given Prob(Error) target
 - There is a fixed SNR, and
 - There is a fundamental tradeoff between
 - SNR, signal, bandwidth, and input capacitance
 - Input capacitance is a proxy for power



References

- [1]
 - Fundamentals of Guided-Wave Optoelectronic Devices, William S.C. Chang
 - Resonant ring modulators
 - "Ultralow power high performance Si photonic transmitter," G. Li et al., OSA OFC/NFOEC (2010)
 - "Ultra-low energy all-CMOS modulator integrated with driver," X. Zheng et al., Optics Express 18 (3), 3059-3070 (2010)
 - "Ultralow power Silicon microdisk modulators and switches," M. Watts et al., IEEE Symposium on Group IV Photonics (2009)
 - QCSE
 - "Device requirements for optical interconnects to silicon chips," D.A.B. Miller, Proc IEEE 97, 1166-1185 (2009)
 - Mach-Zehnder
 - "A Fully Integrated 4 10-Gb/s DWDM Optoelectronic Transceiver Implemented in a Standard 0.13um CMOS SOI Technology," A. Narasimha et al., IEEE JSSC 42 (12), 2736-2744 (2007)
- [2]
 - Design of High-Speed Optical Interconnect Transceivers, Sam Palermo, Ph.D. Thesis



References, con't

- [3]
 - "Wide temperature range operation of micro-meter scale silicon electro-optic modulators," Manipatruni et al, Optics Letters 33:19, 2185-7 (2008)
- [4]
 - "CMOS-integrated 40GHz Germanium waveguide photodetector for on-chip optical interconnects," S. Assefa, OSA OFC/NFOEC (2009)
- [5]
 - Microelectronic Circuits, Sedra and Smith
- [6]
 - "State of the art Si-based receiver solutions for short reach applications," M. Morse et al., OSA OFC/NFOEC (2009)
 - "A sub-picojoule-per-bit CMOS photonic receiver for densely integrated systems," X. Zheng et al., Optics Express 18:1, 204-211 (2010)



References, con't

- [7]
 - “Progress in low-power switched optical interconnects,” Krishnamoorthy et al., IEEE Journal for Selected Topics in Quantum Electronics, in press for March 2011.
- [8]
 - “Computer systems based on silicon photonic interconnects,” Krishnamoorthy et al., Proc. IEEE 97:7, 1337-1361 (2009)
 - “Circuits for silicon photonics on a macrochip,” Ho et al., IEEE Asian Solid State Circuits Conference, 182-183 (2009)
 - “Optical interconnect for high-end computer systems,” Ho et al., IEEE Design and Test of Computers 27:4, 10-19 (2010)