50.4 Power-Centric Design of High-Speed I/Os

Hamid Hatamkhani[3], Frank Lambrecht[2],
Vladimir Stojanovic[1],
and Chih-Kong Ken Yang[3]

[1] MIT, [2] Rambus Inc., [3] UCLA

I/O Bandwidth Requirements: Computers

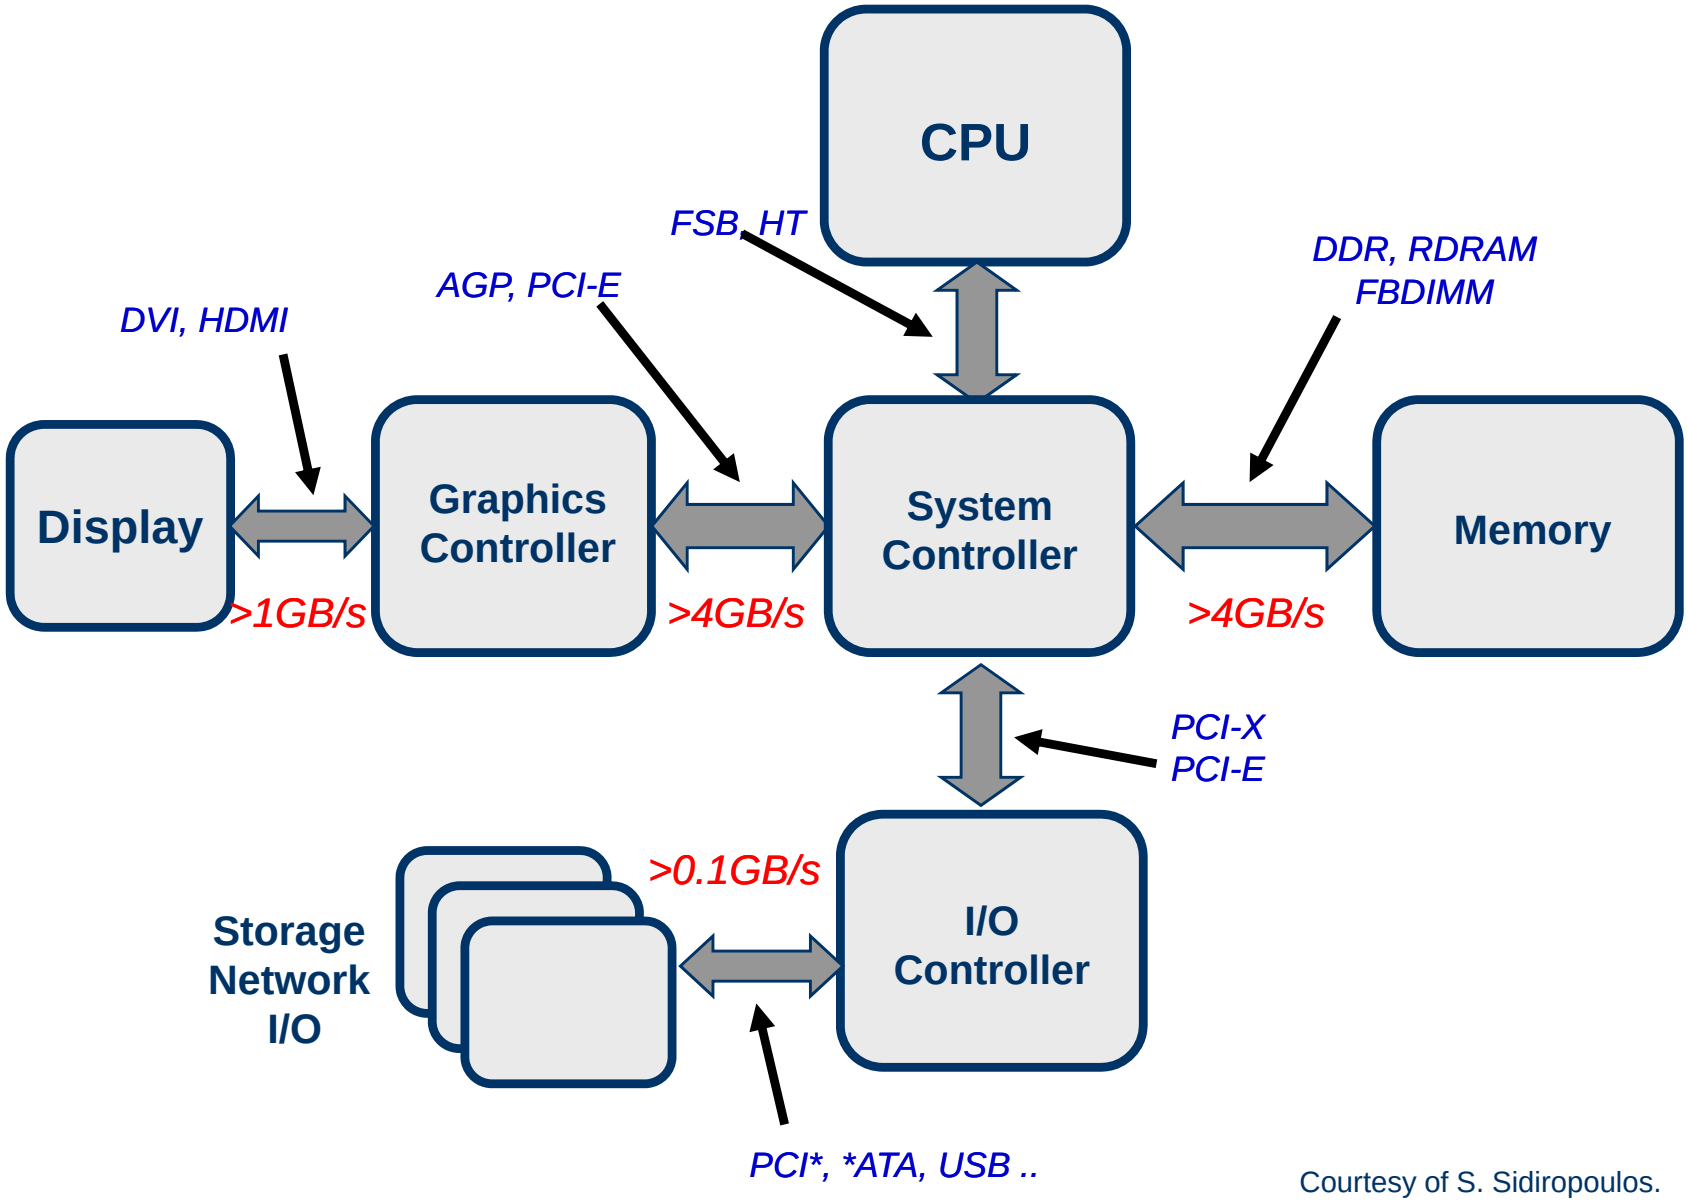

Courtesy of S. Sidiropoulos.

I/O Bandwidth Requirements: Network Systems

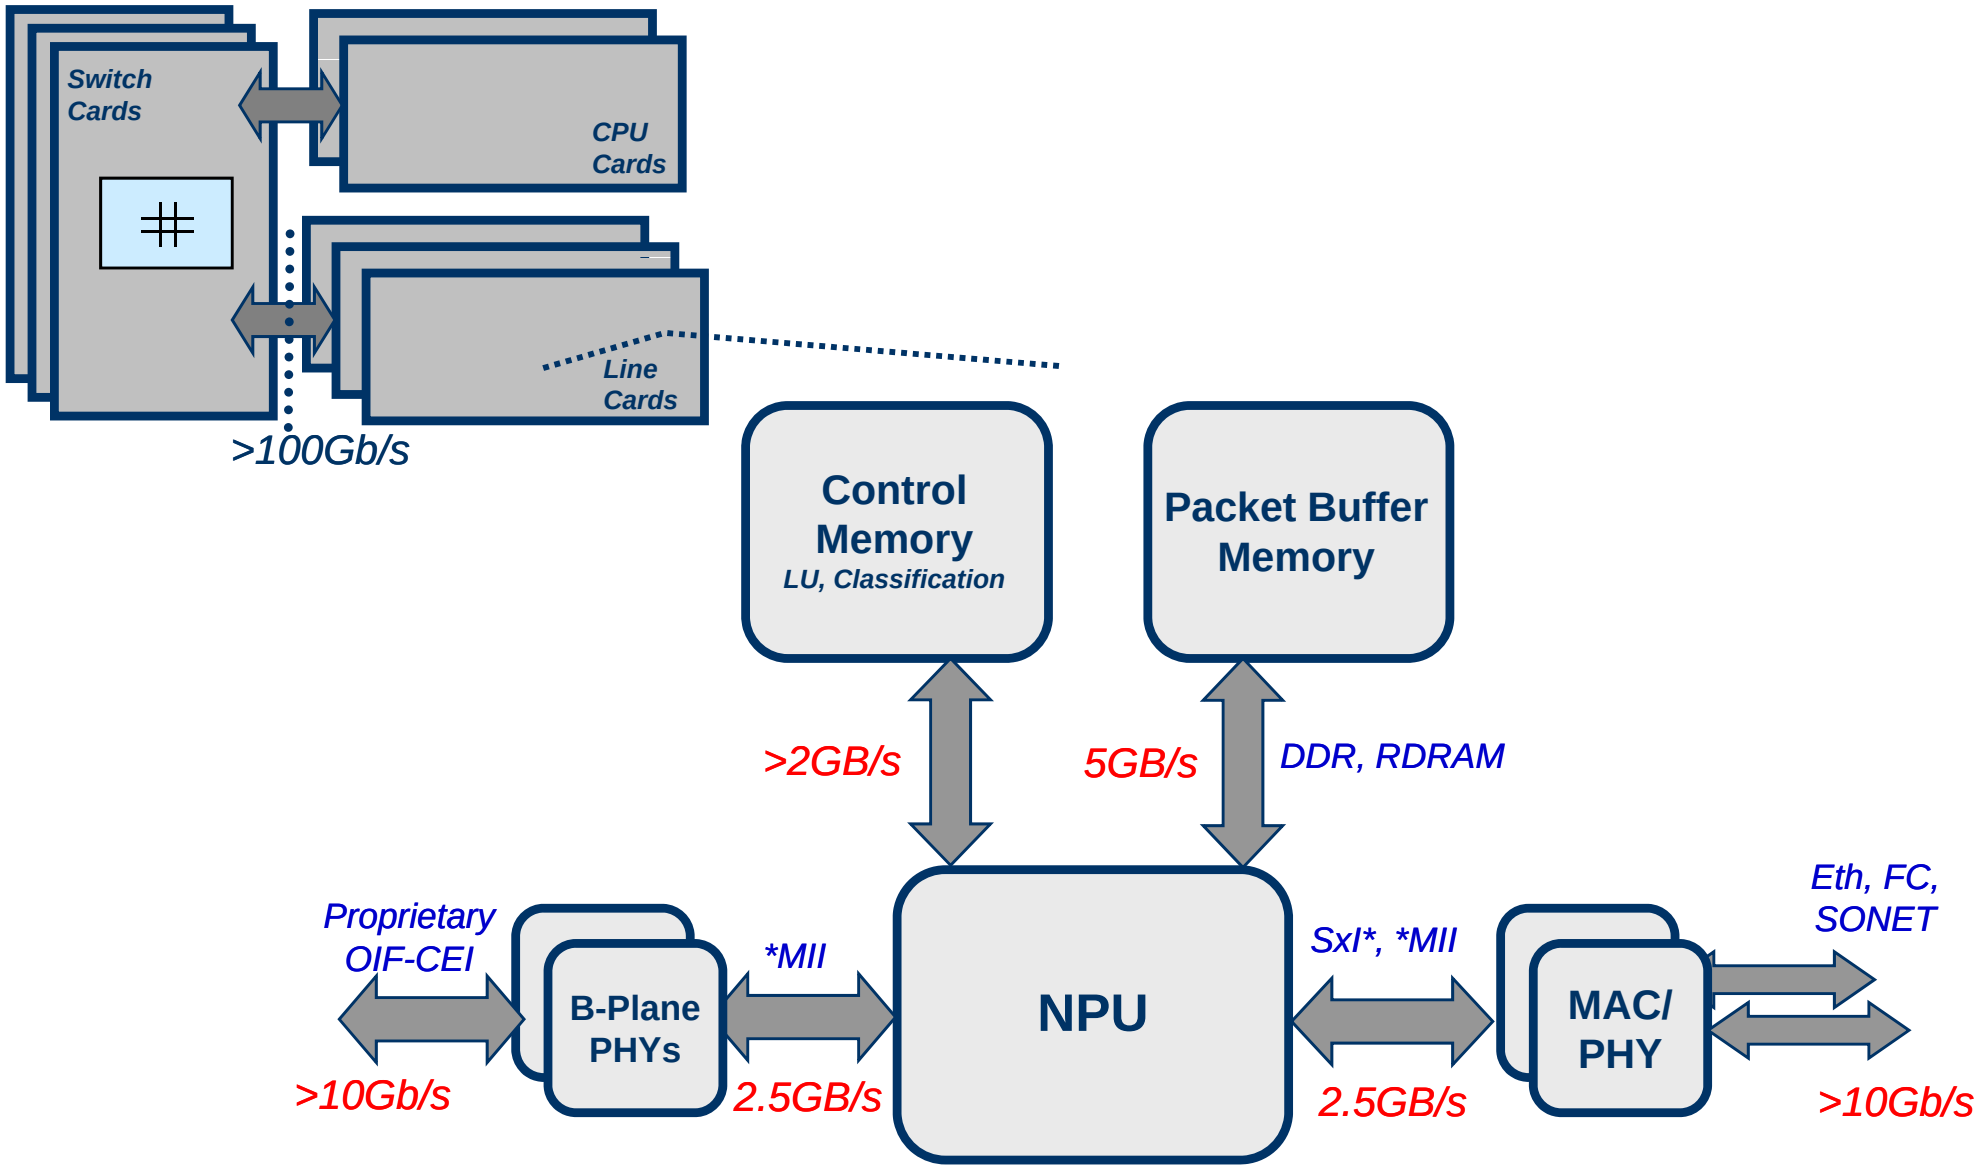

Courtesy of S. Sidiropoulos.

ITRS Bandwidth Projection

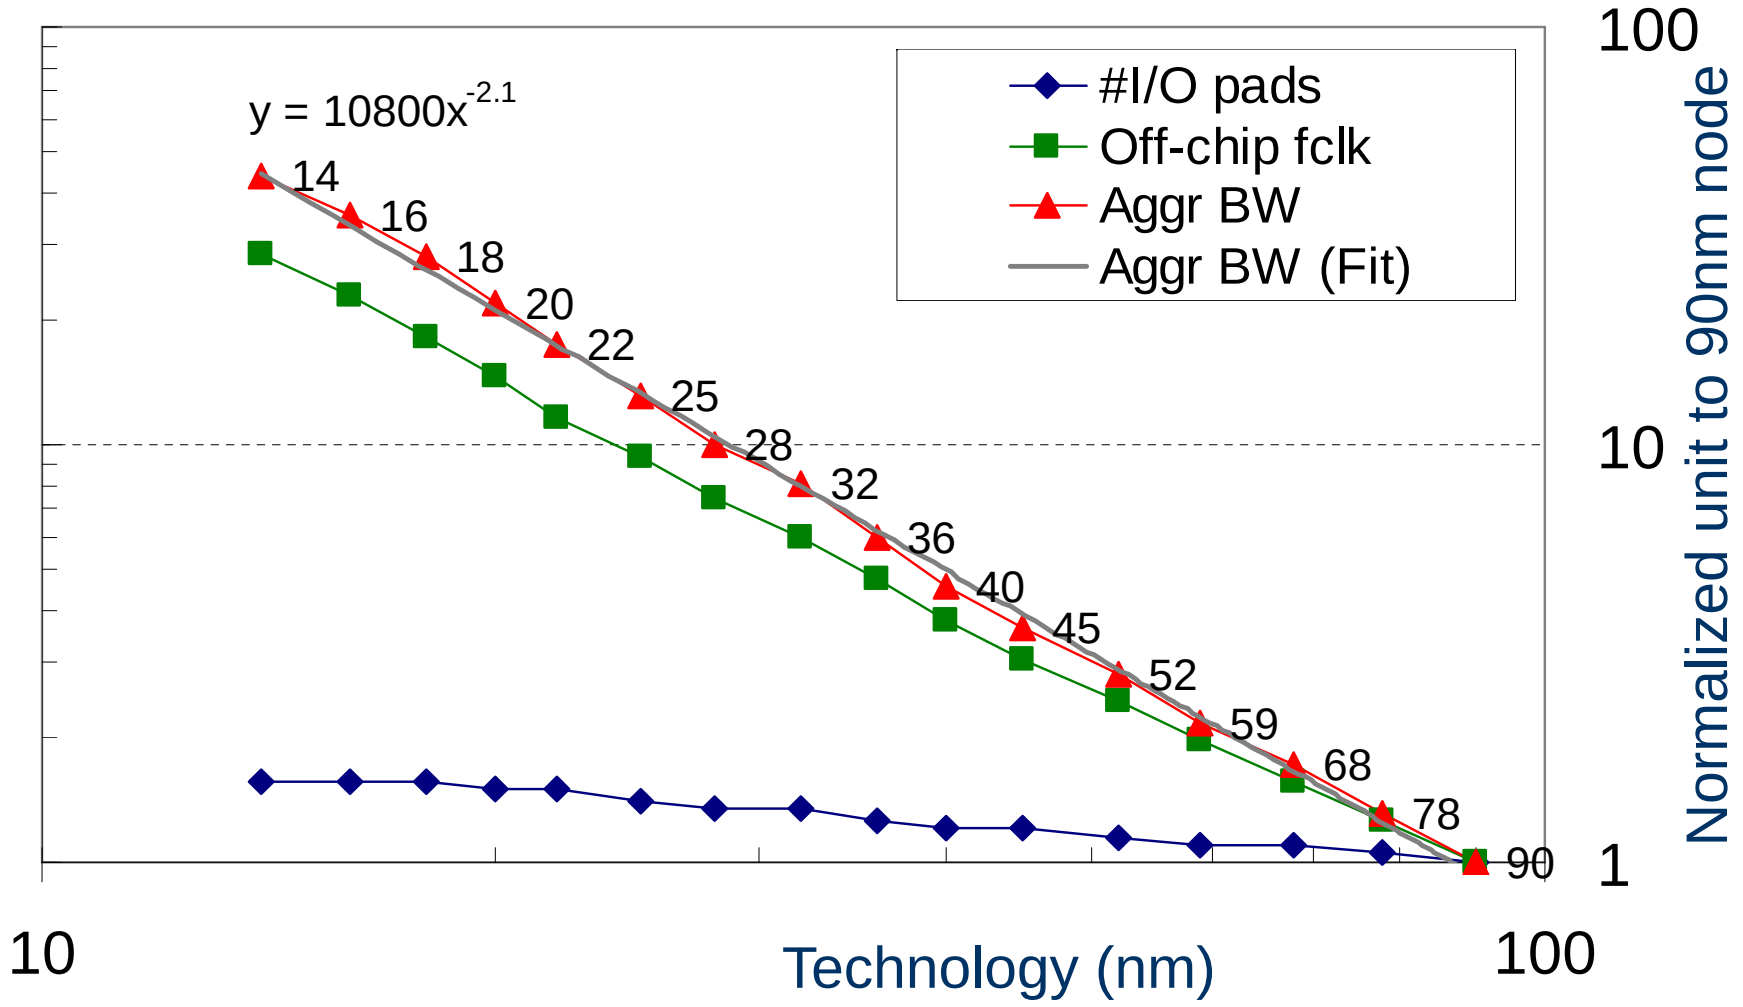

Power Survey

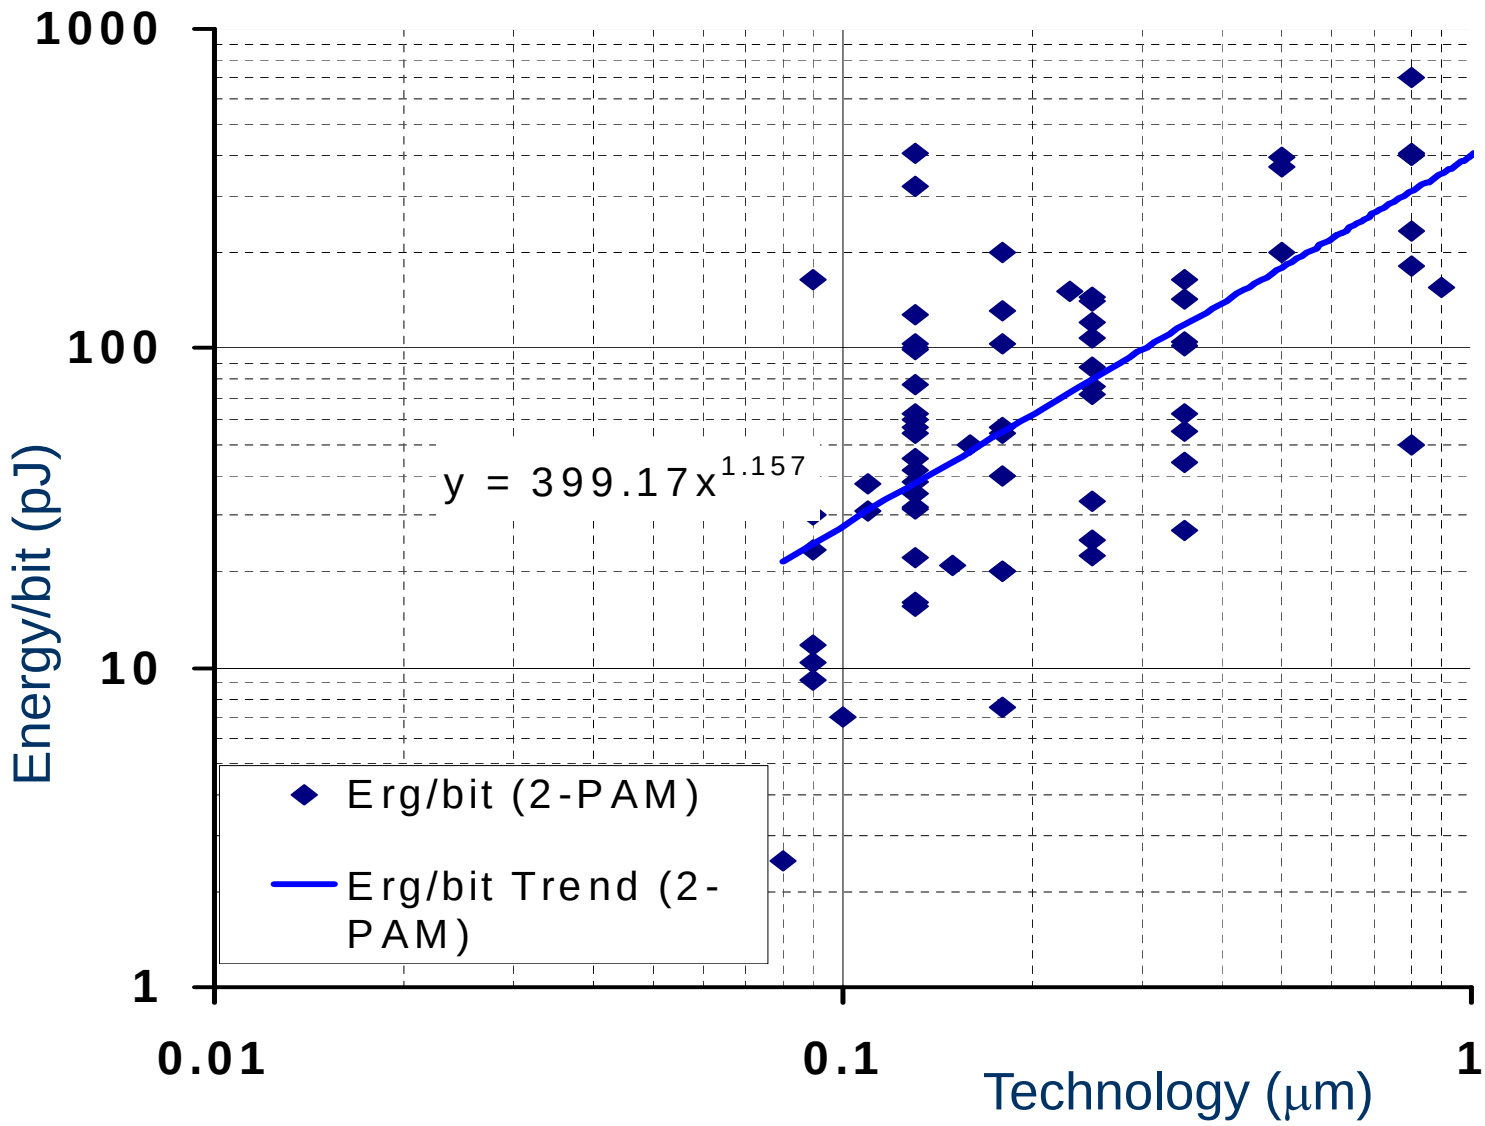

Outline

- Introduction to an I/O System
- Power-Centric Design Methodology
- Examples of Design Tradeoffs
- Conclusion

I/O Building Blocks

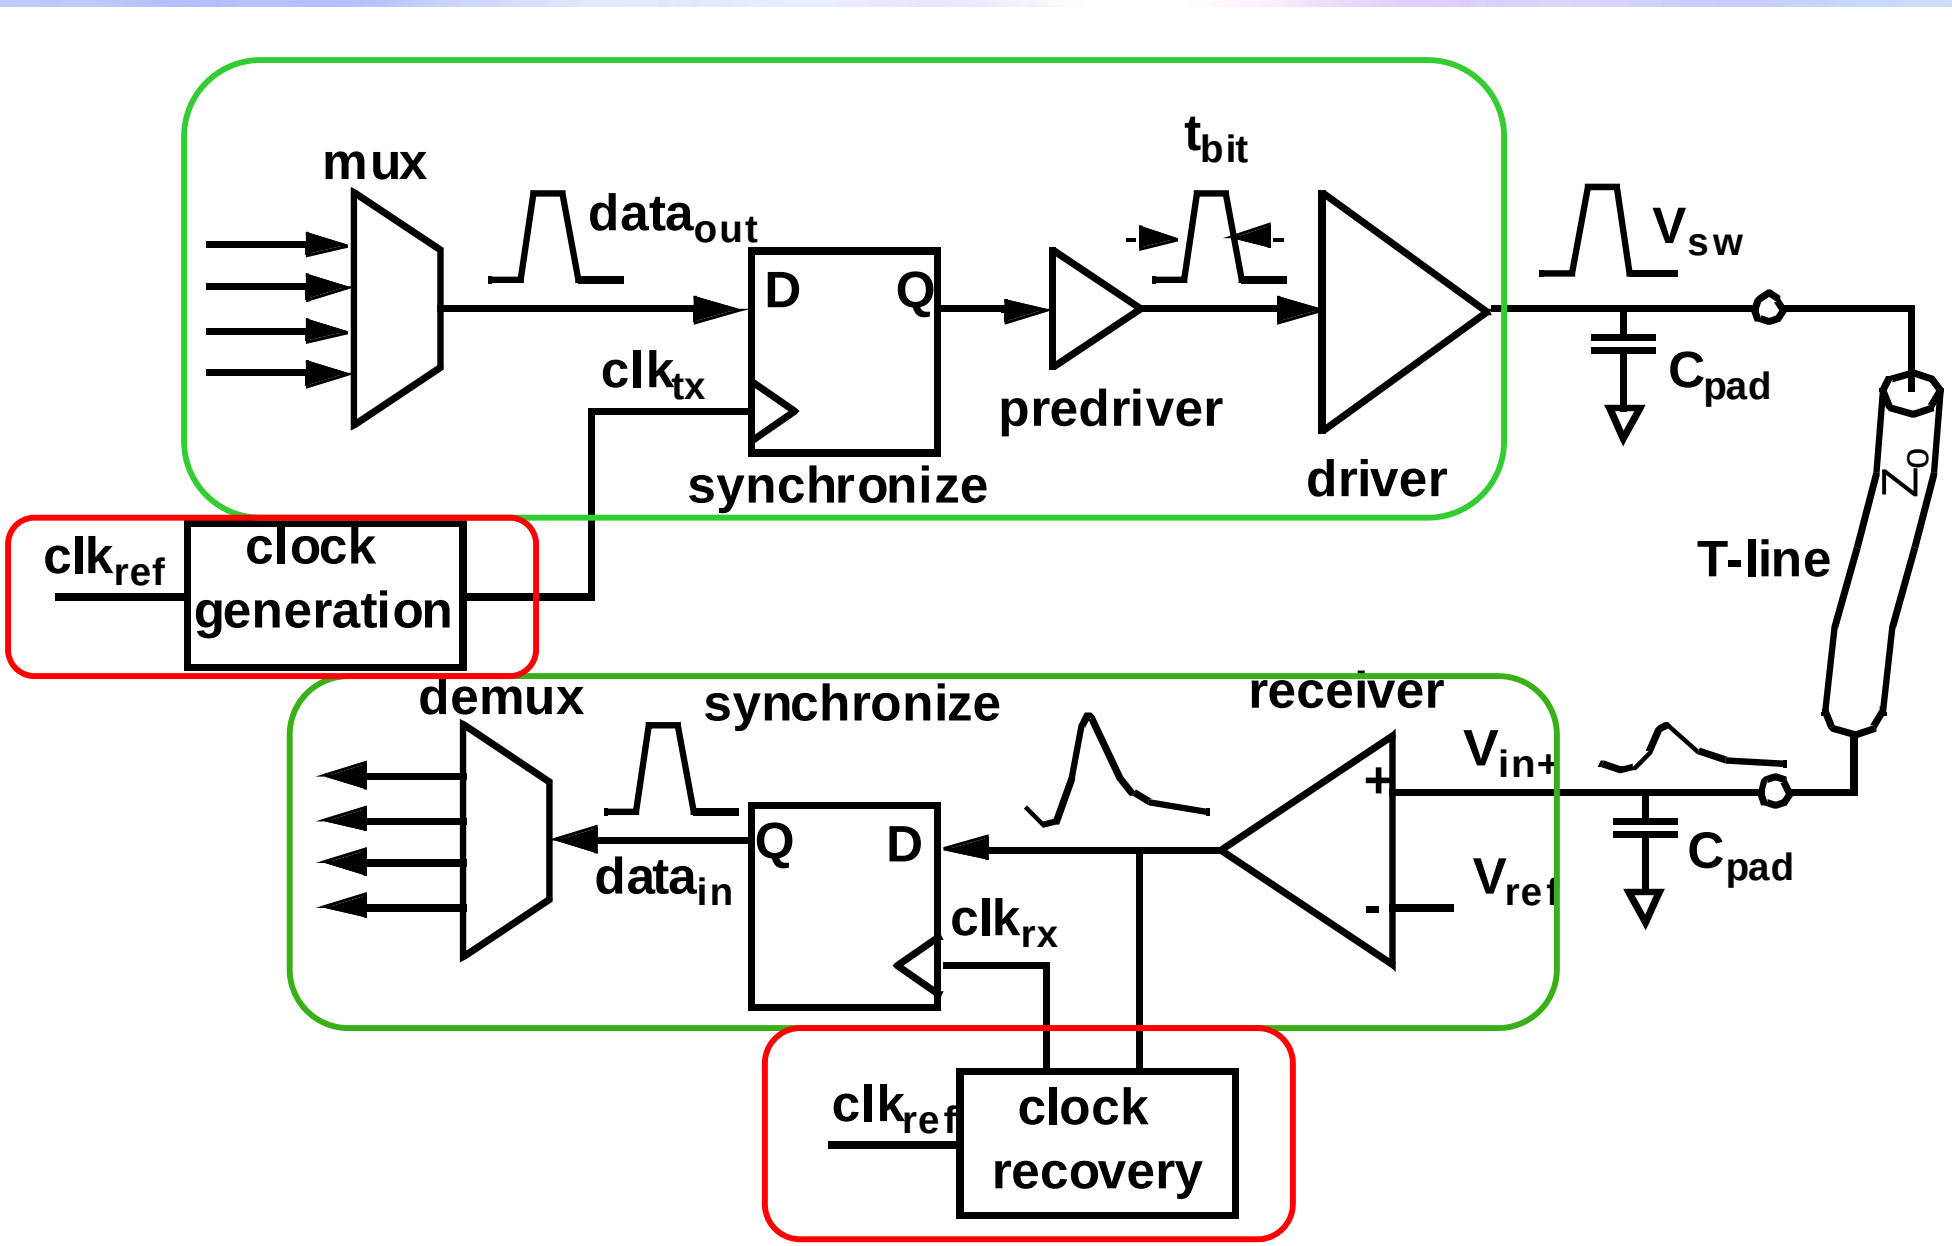

Circuit Design Goals

- Primary issue is noise and interference.
- Minimize circuit noise
  - Timing noise
  - Voltage noise
- Compensate (filtering) for channel and noise.

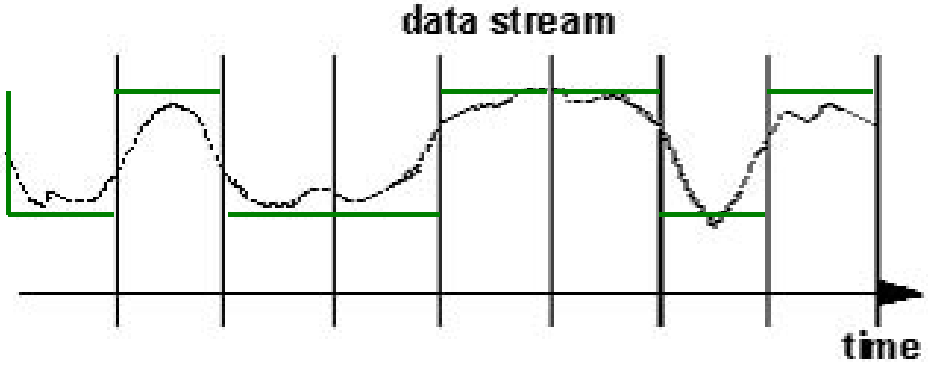

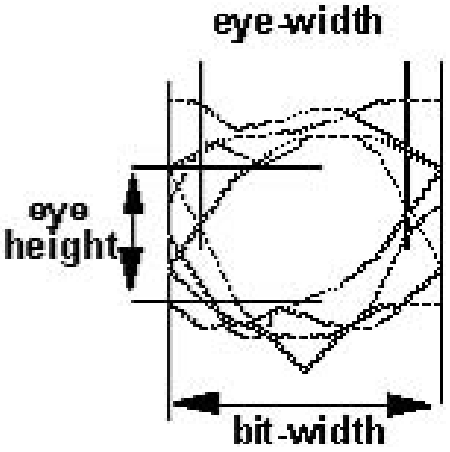

Transmission Channel

- Connection between chips
  - Many components: bondwire, package, PCB, connectors, cables
  - Backplane example:

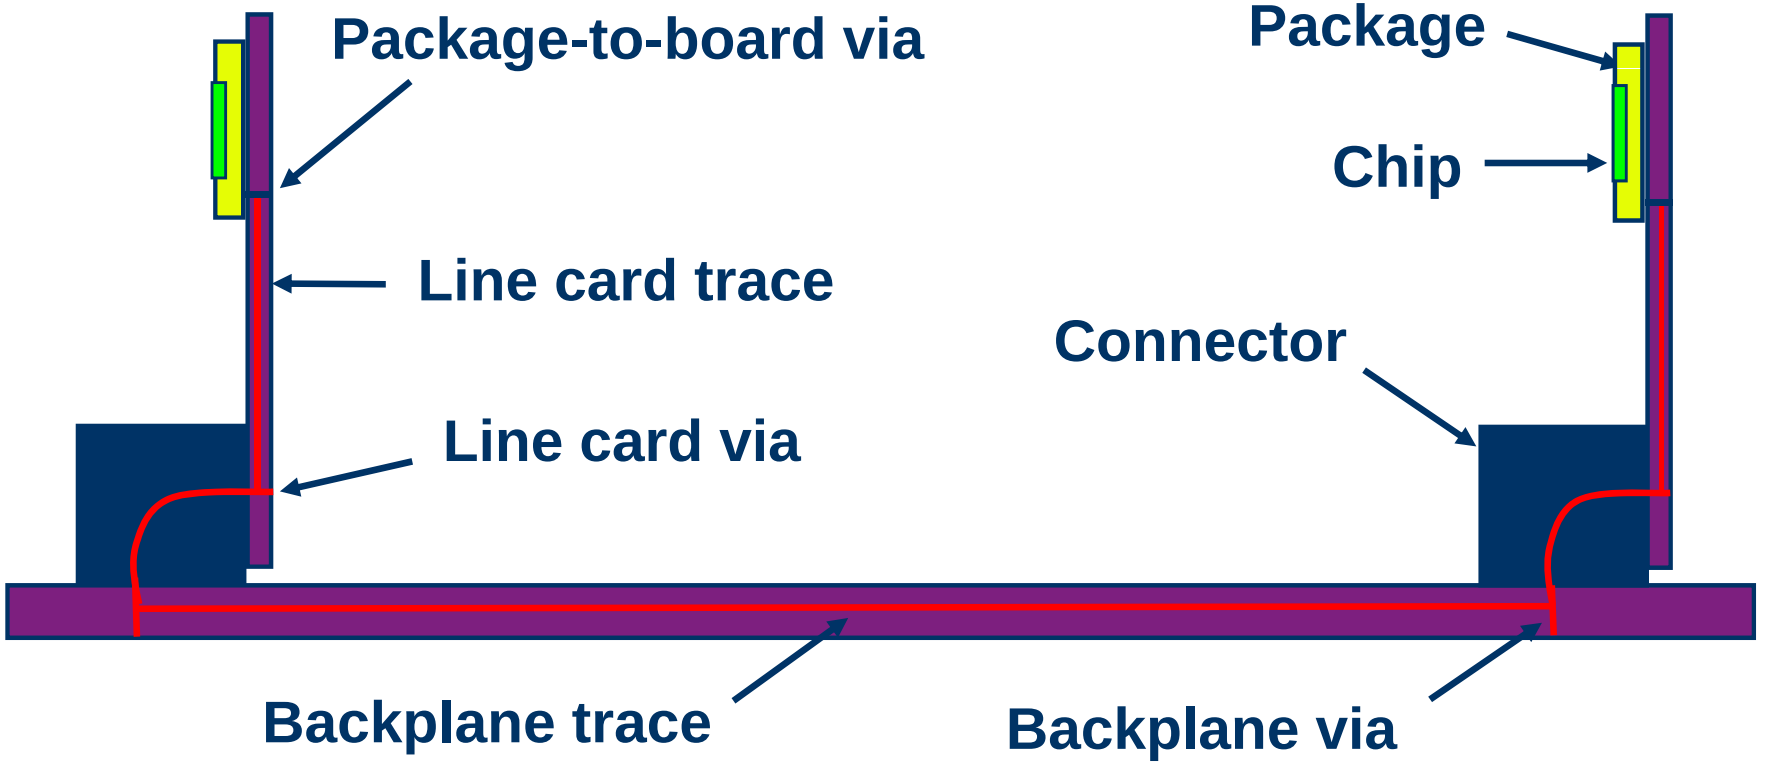

Courtesy of J. Zerbe.

Channel Response

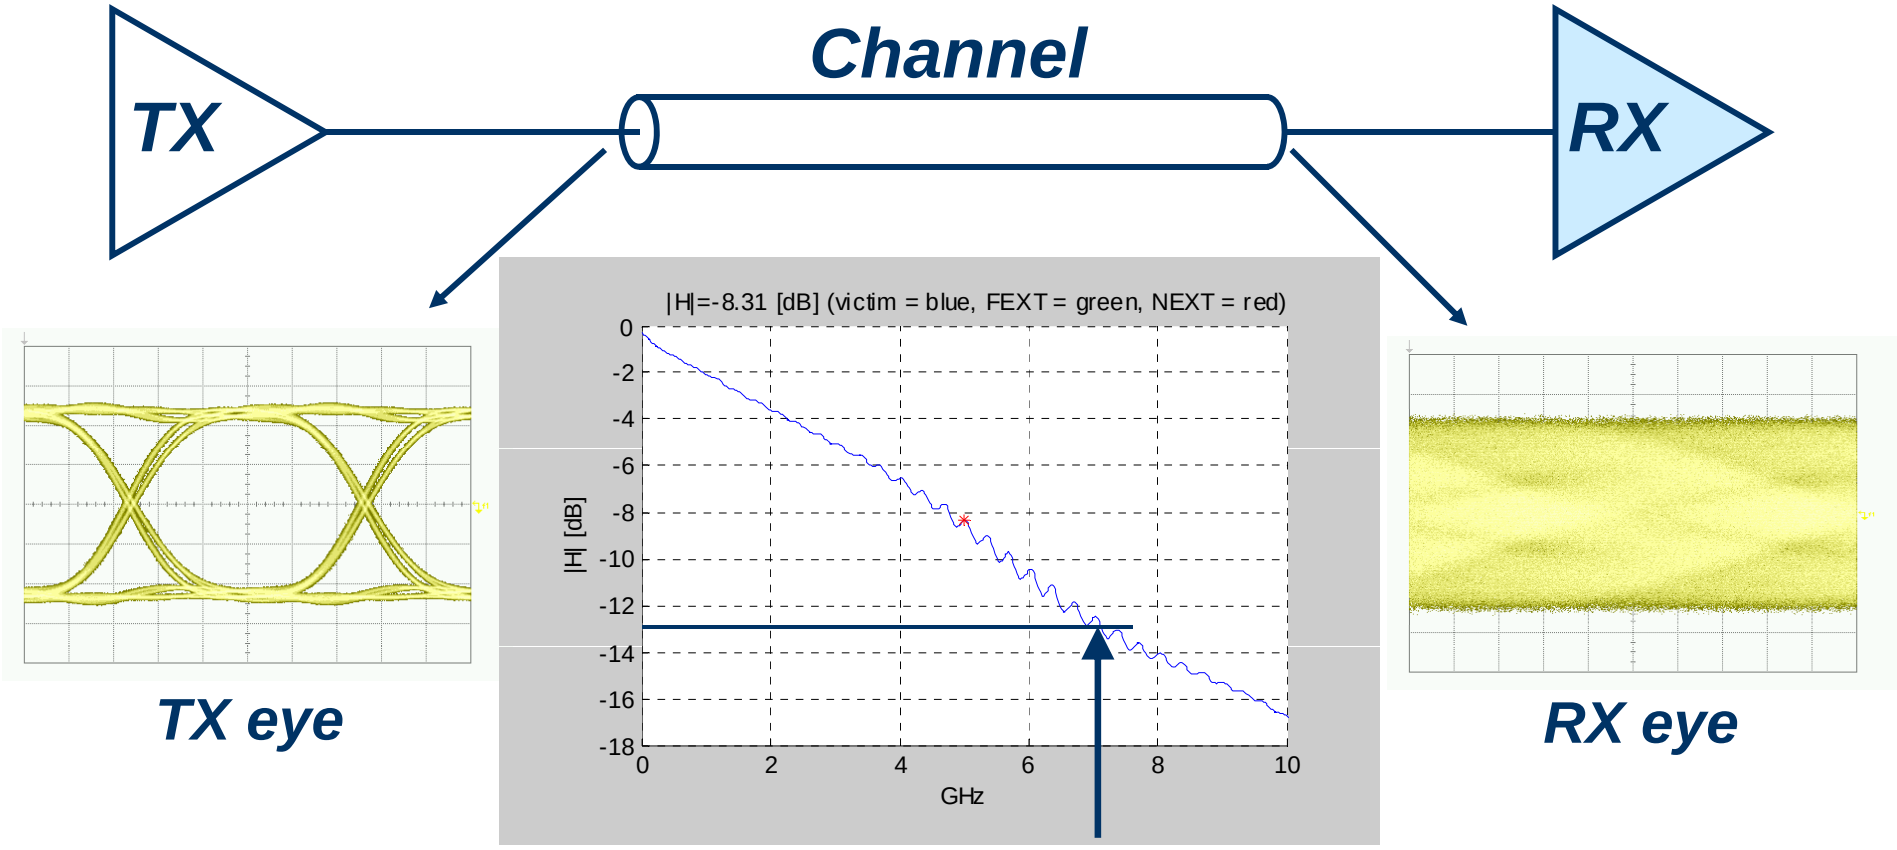

- Pulse spreading closes the data eye
  - Inter-symbol interference (ISI)

Compensating for ISI

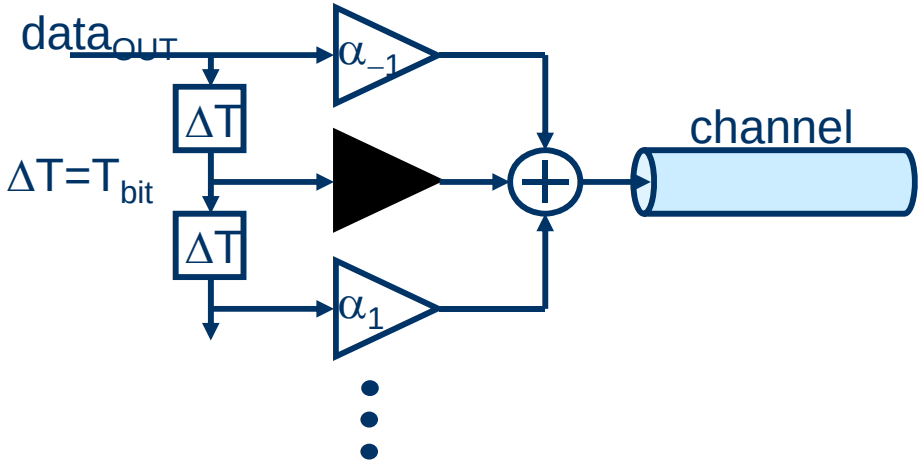

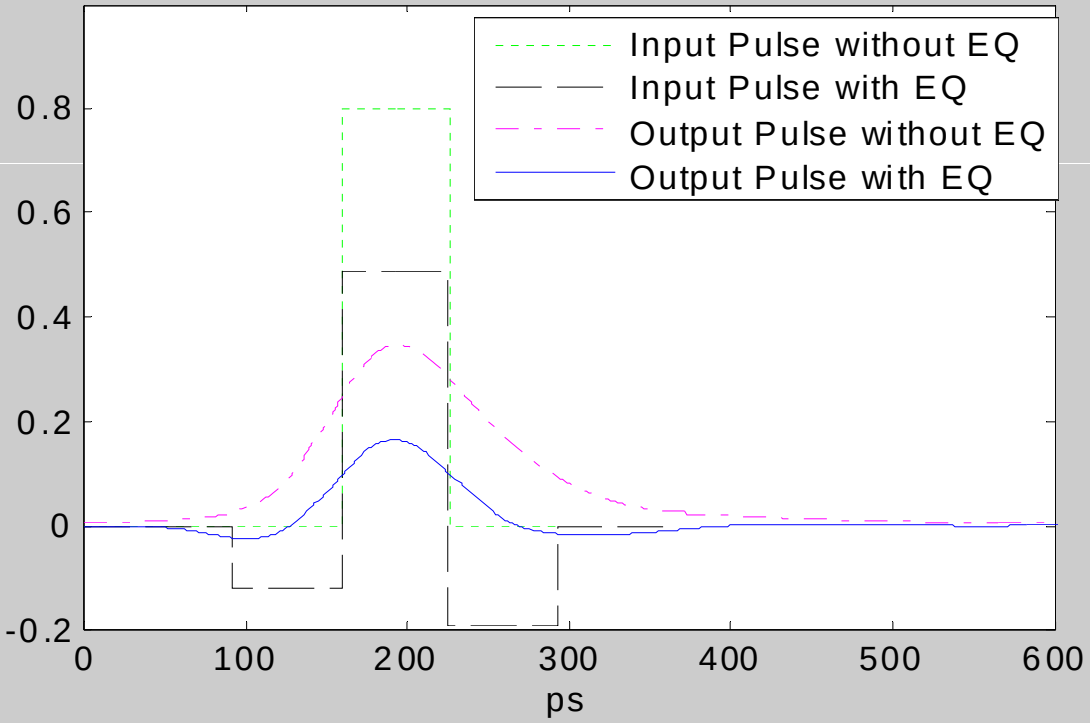

- Filter to flatten the frequency response.
  - TX = pre-distort the transmitted waveform.
  - RX = equalize with a high-pass filter.

Eye Mask Specification

- Example from Infiniband
  - UI = unit interval, normalized T_{bit}.
  - Deterministic jitter (DJ), and random jitter (RJ).
- Becoming more sophisticated with more equalization.

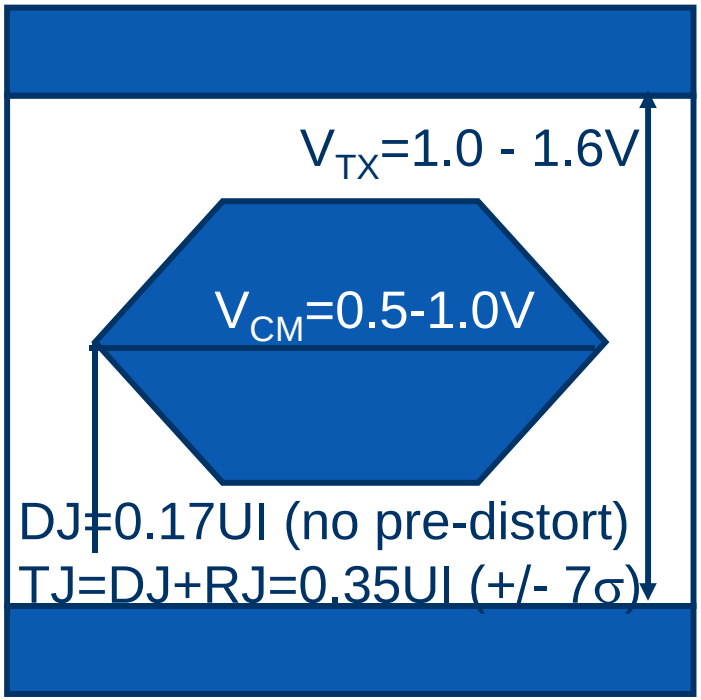

Transmit Side

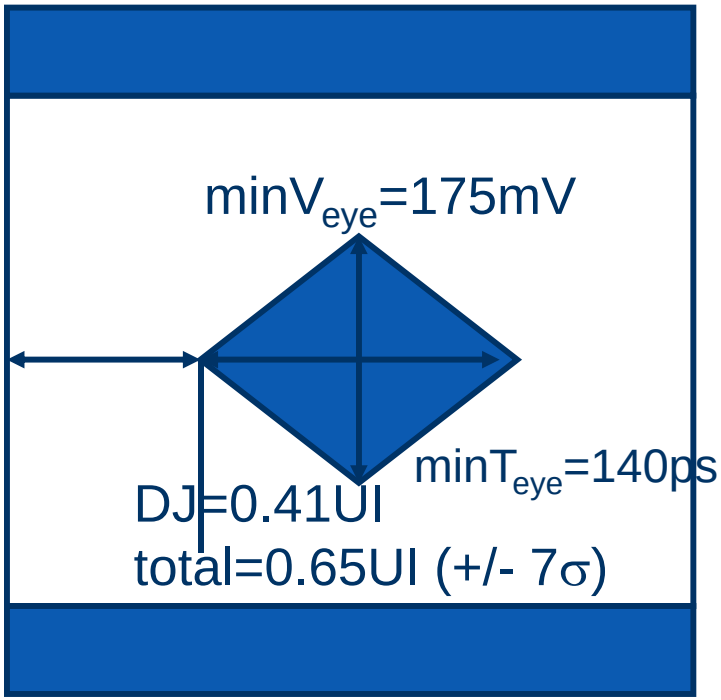

Receive Side

Receive Margin and BER Specification

- Plot the BER as the receiver sampling point (time+voltage) is swept
  - Specify amount of receiver margin before BER>10^{-11}

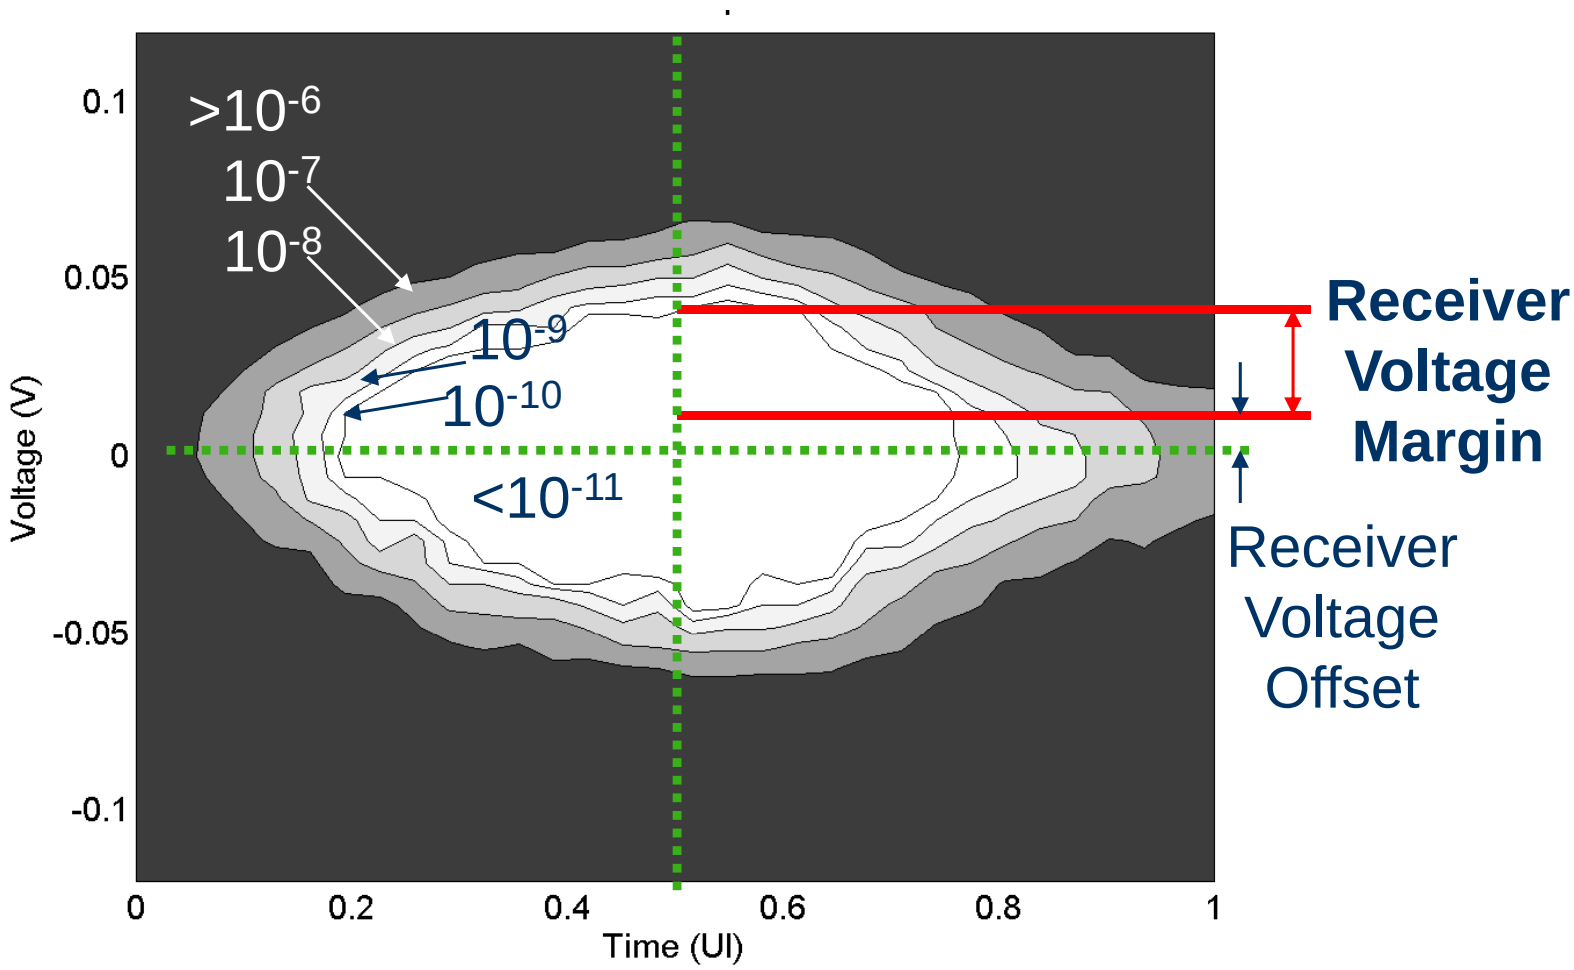

Outline

- Introduction to an I/O System
- Power-Centric Design Methodology
- Examples of Design Tradeoffs
- Conclusion

I/O Power Dissipation

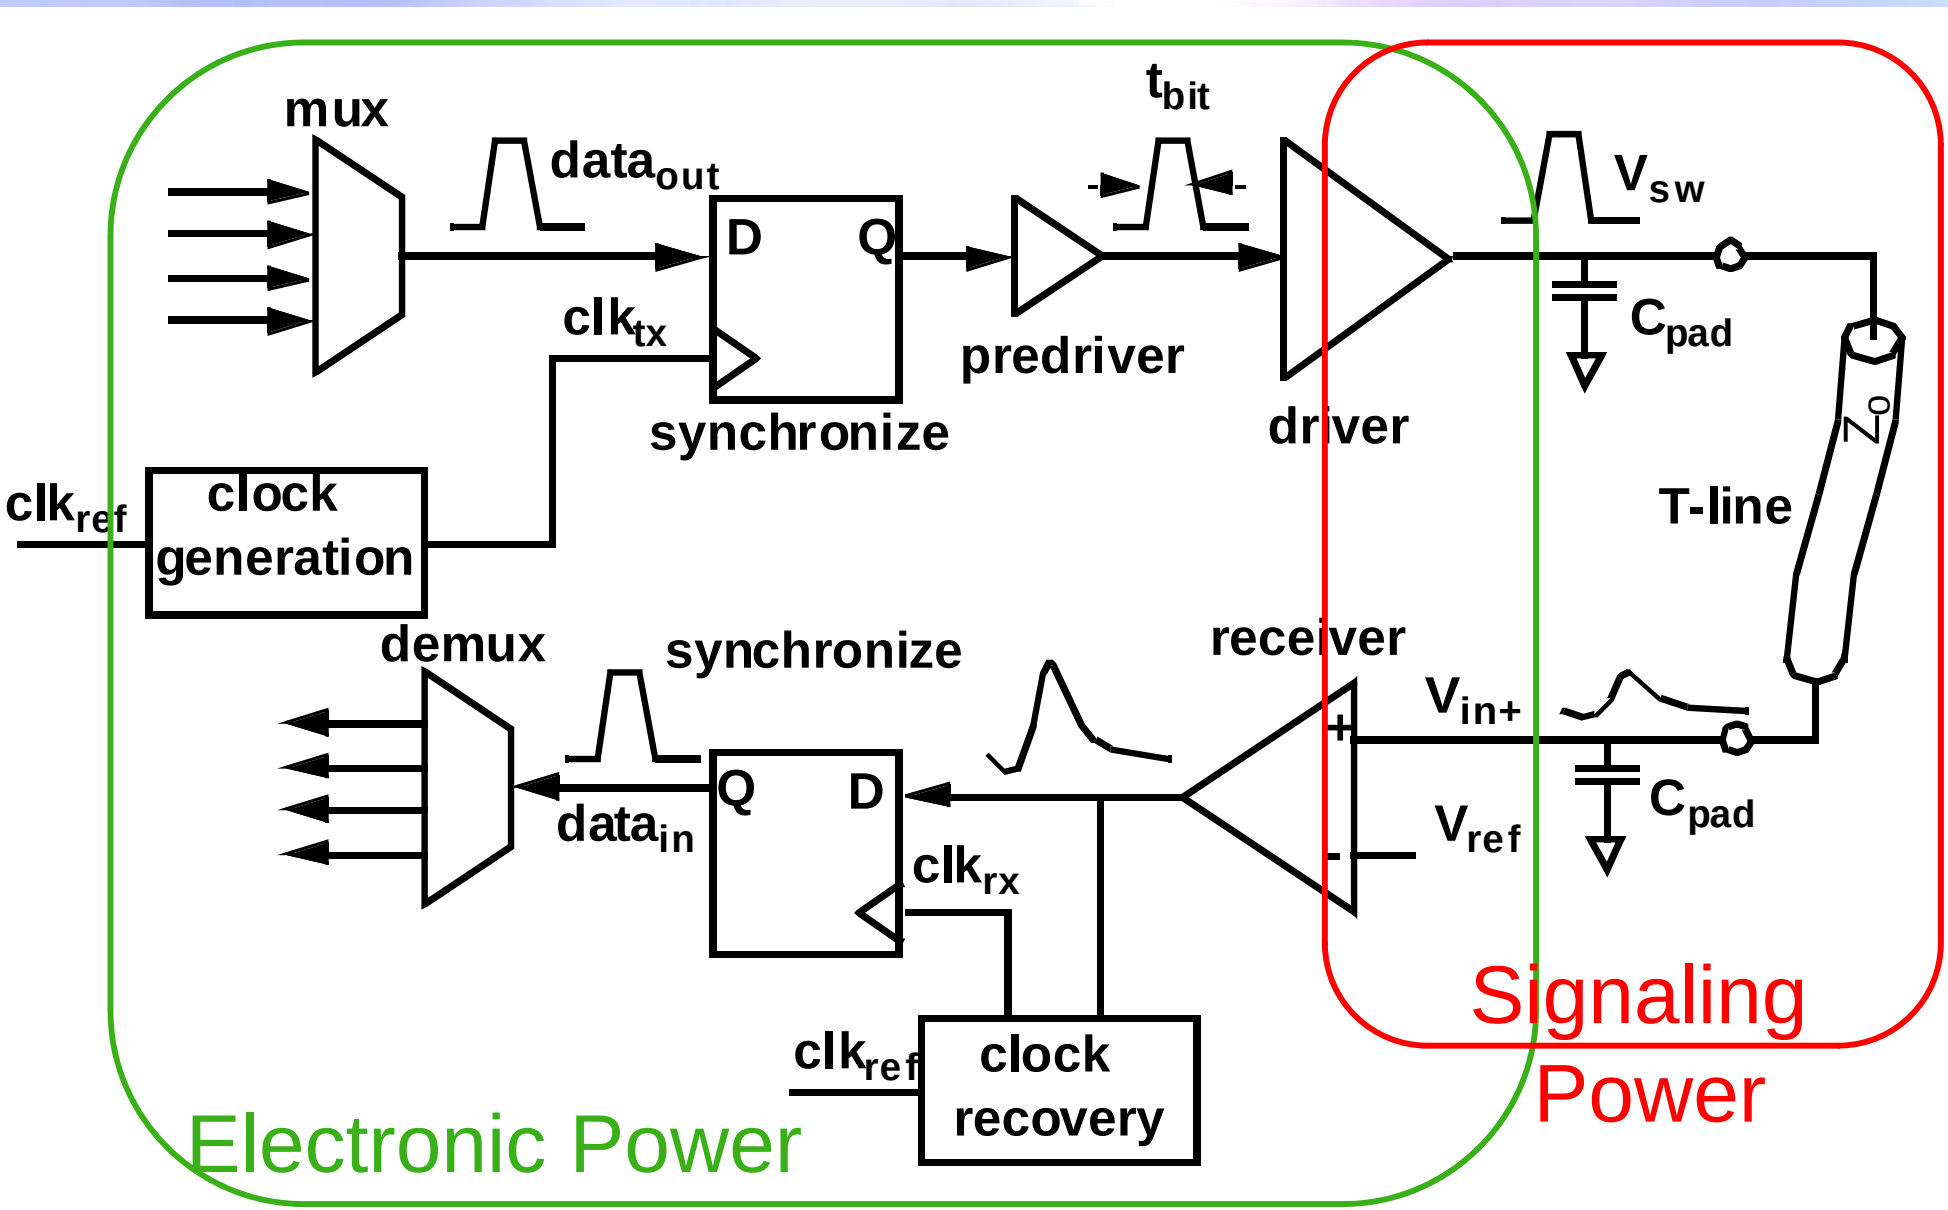

Signaling Power

- Transmitting onto a 50Ω channel
  - Depends on the swing
  - Depends on the common-mode voltage
- Energy/bit improves with higher data rate.

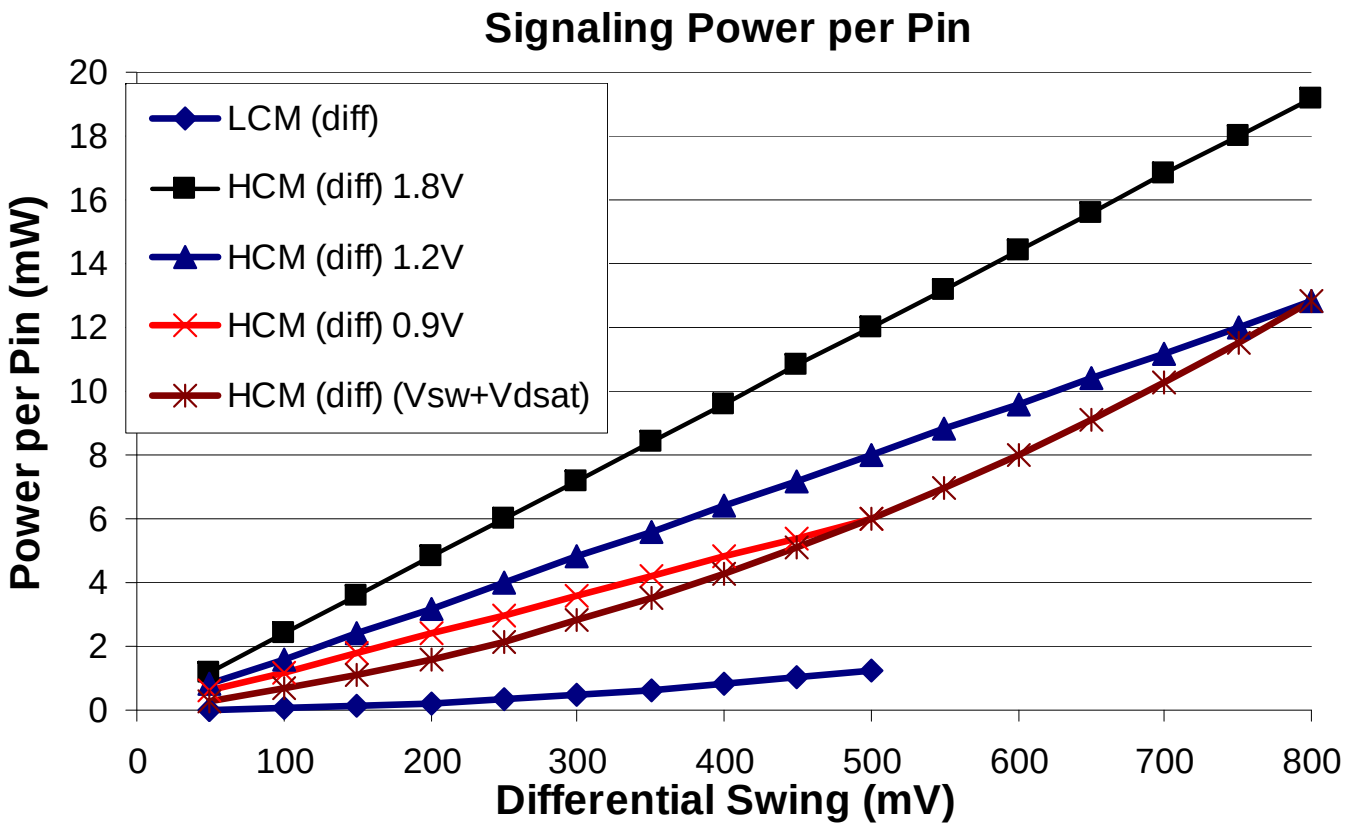

Low Electronic Power Design

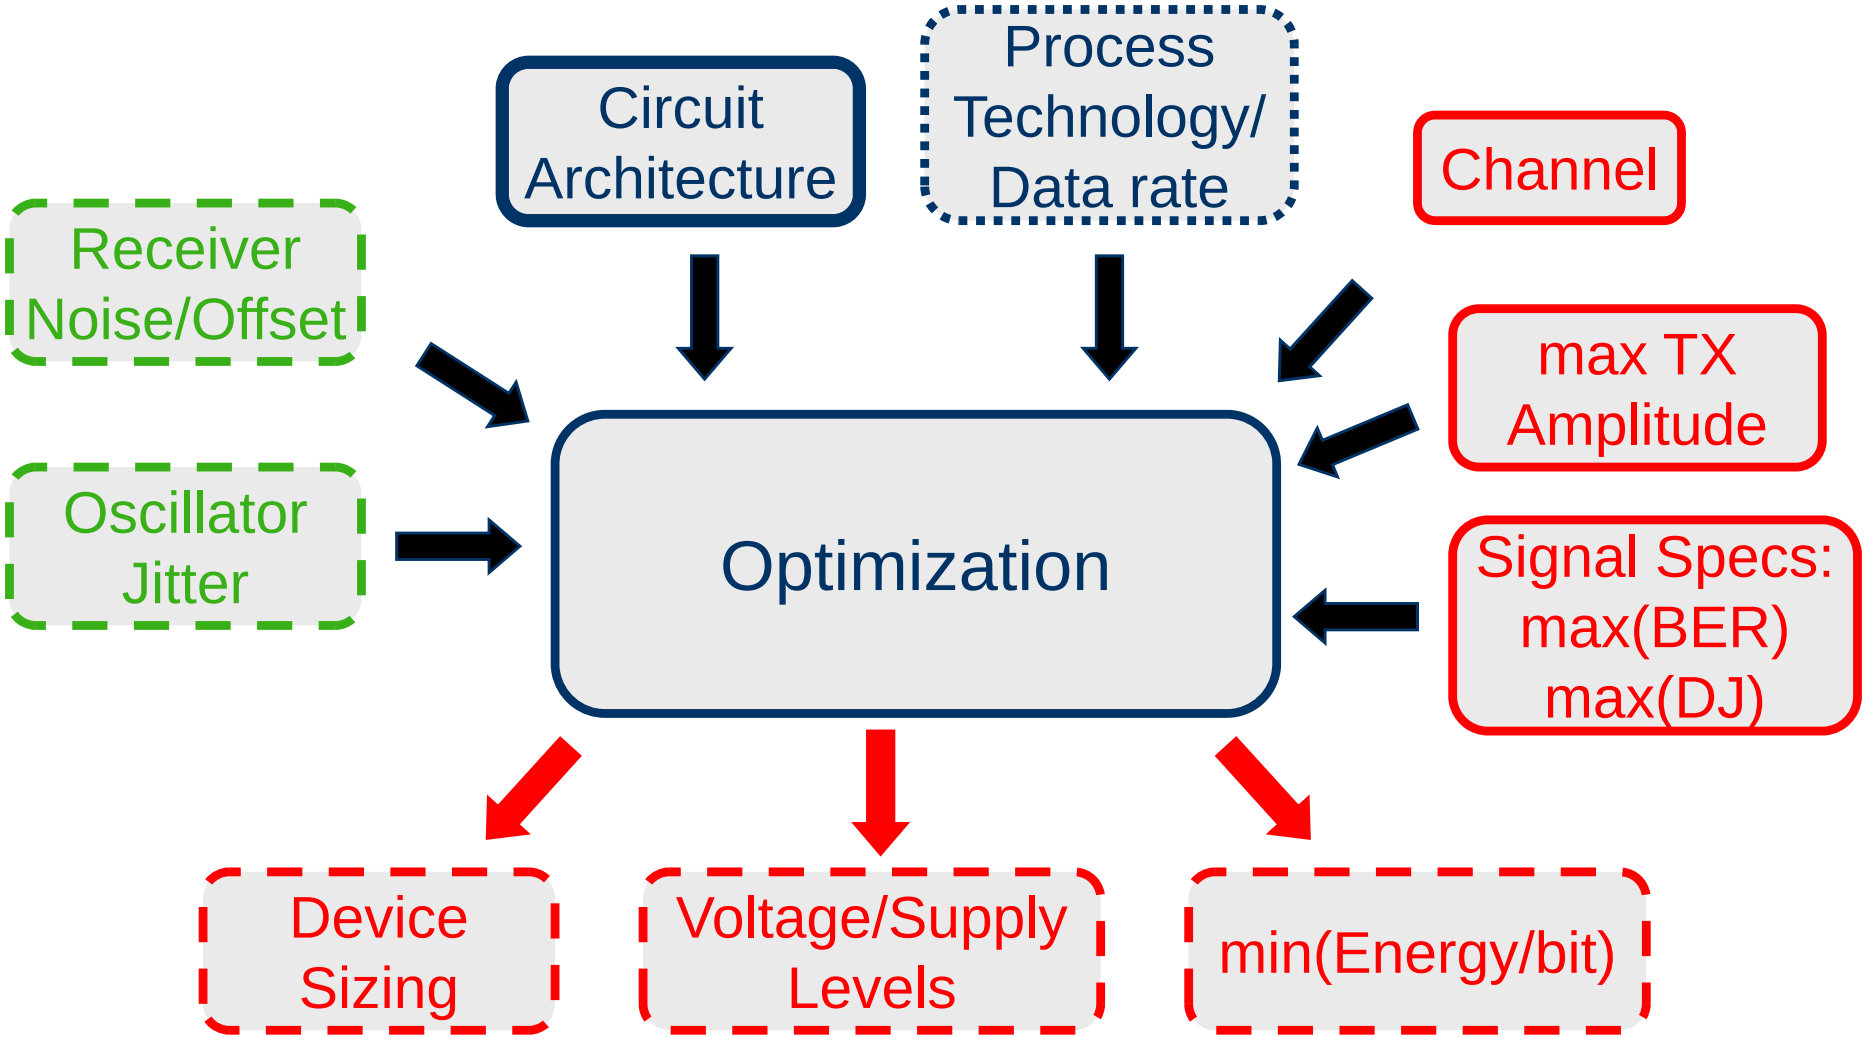

Example: Circuit Architecture

- Varying common-mode levels require different driver designs
  - Designs are mixed-signal.
- Logic family depends on speed (i.e. CML)

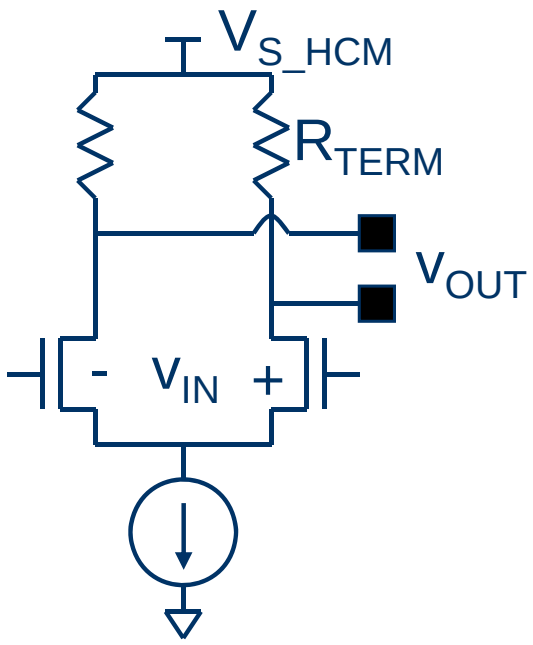

High-CM Driver

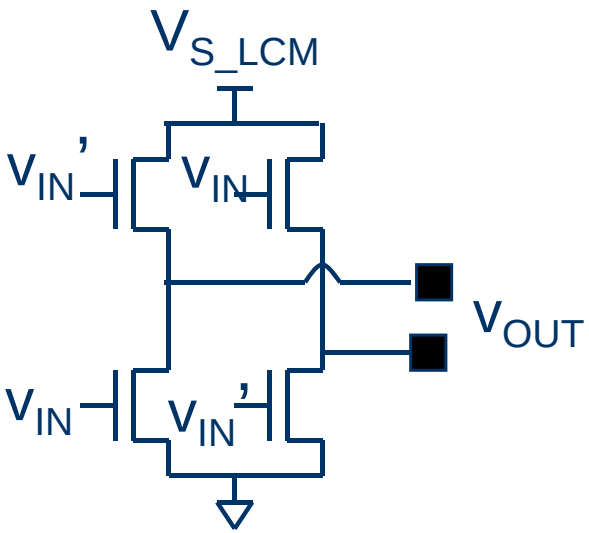

Low-CM Driver

Example: Logic Architecture

- Position of multiplexing
  - Closer to the output allows higher data rates but at higher power dissipation.

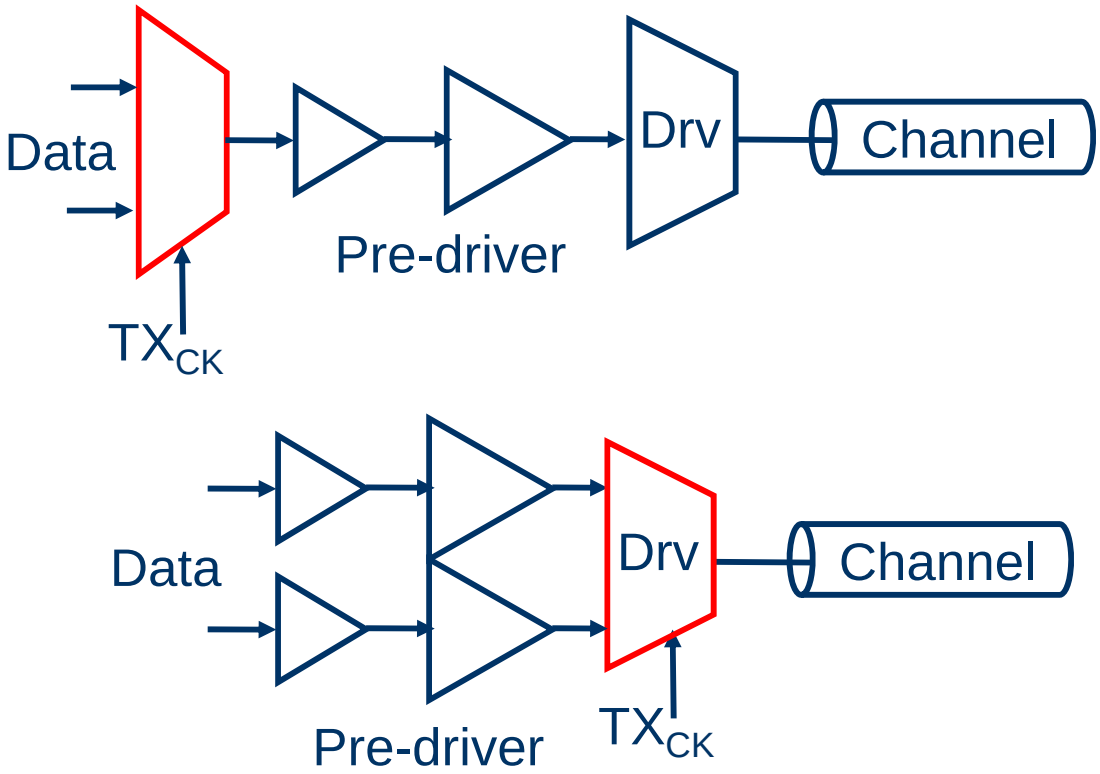

Example: Compensating for Interference

- Employing slew-rate control
  - Filters out high-frequency information to minimize noise coupling
- Pre-distort to compensate ISI
  - i.e. FIR filter: $H(z) = \alpha_{-1}z + 1 + \alpha_1 z^{-1} + \alpha_2 z^{-2} + \ldots$

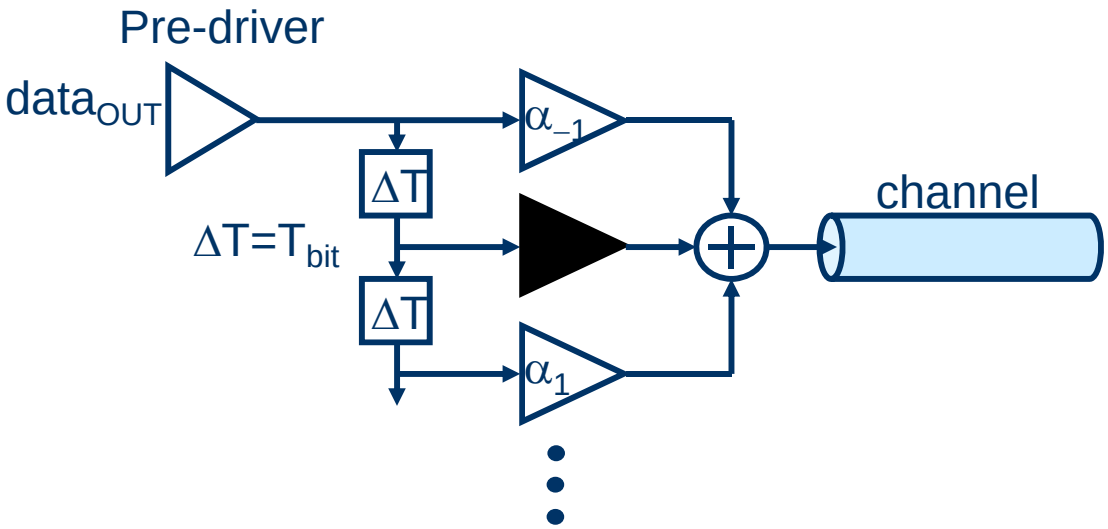

Modified Circuit Analysis/Optimization

- Optimize circuits with added constraints
  - Slew-rates
  - Timing variations
  - Noise
  - Specialized ”analog” circuit structures
- Example: Deterministic Jitter (DJ)

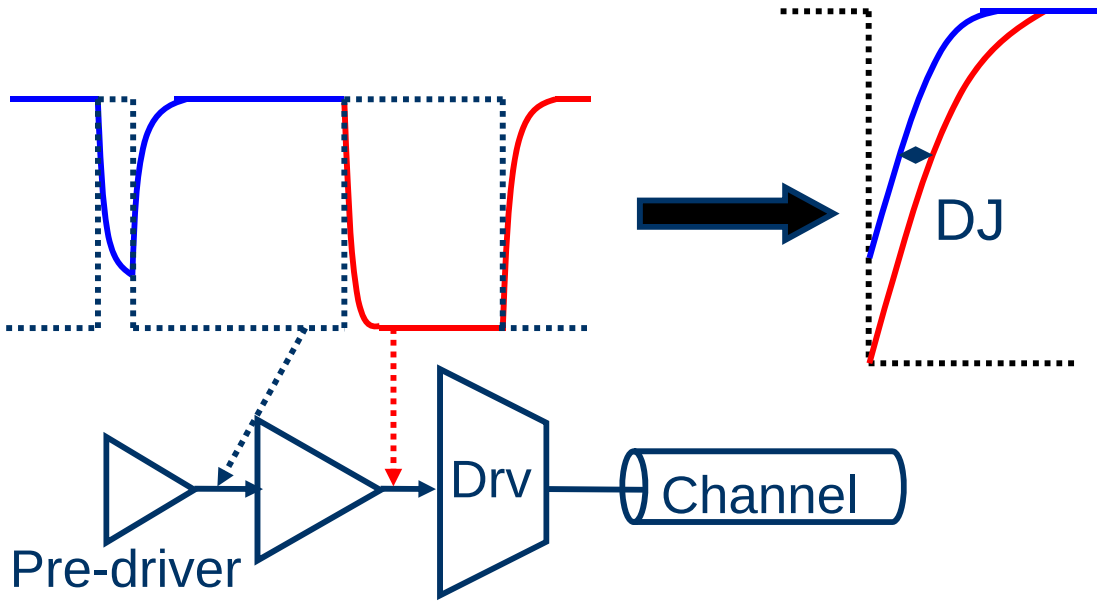

Evaluating the BER

- Combine circuit optimizer to an oracle of the channel
  - Statistical channel model – LinkLab (Rambus)

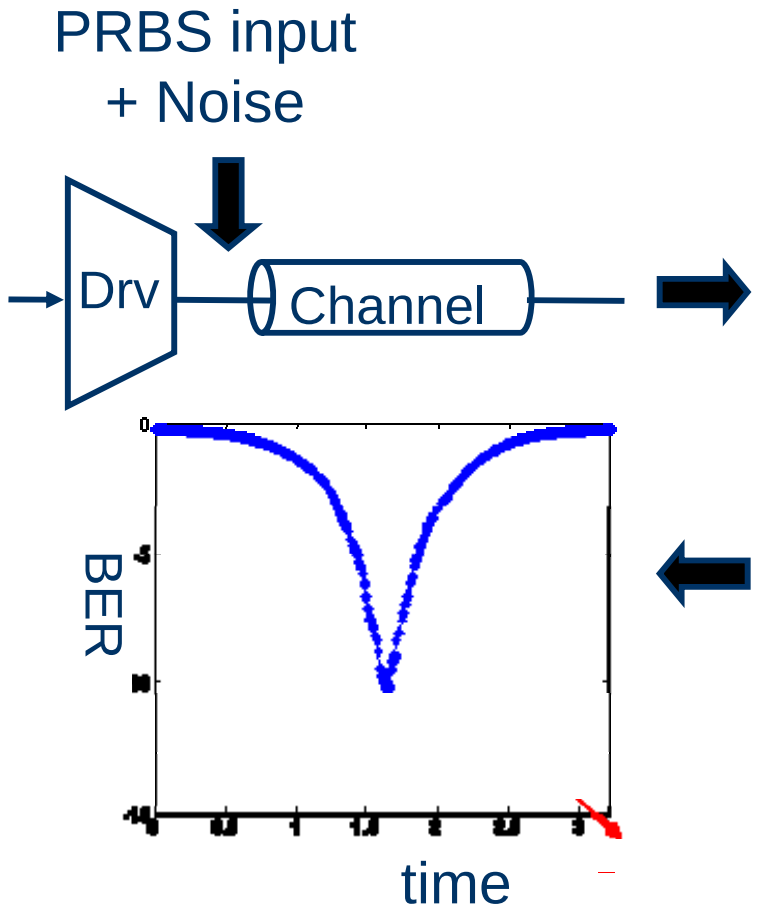

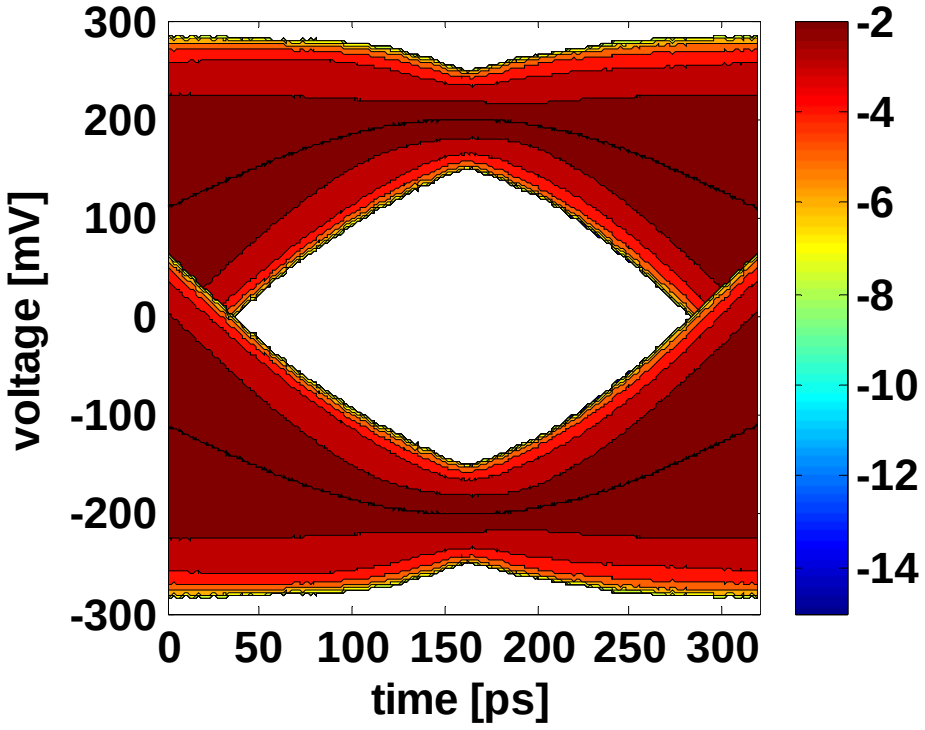

LinkLab Analysis

- Decomposes the signal into pulse stream and noise.
  - Convolve with channel
- Add timing recovery probabilities

$$x_k = x_k^{ISI} + x_k^{jitTX} + x_k^{jitRX}$$

$$\begin{cases} x_k^{jitRx} = \varepsilon_k^{RX} \sum_{n=-sbS}^{sbE} b_{k-n}\left(h_n - h_{n-1}\right) \\ x_k^{ISI} = \sum_{n=-sbS}^{sbE} b_{k-n} p_n \\ x_k^{jitTx} = \sum_{n=-sbS}^{sbE} b_{k-n}\left(h_{n-1}\varepsilon_{k-n+1}^{TX} - h_n \varepsilon_{k-n}^{TX}\right) \end{cases}$$

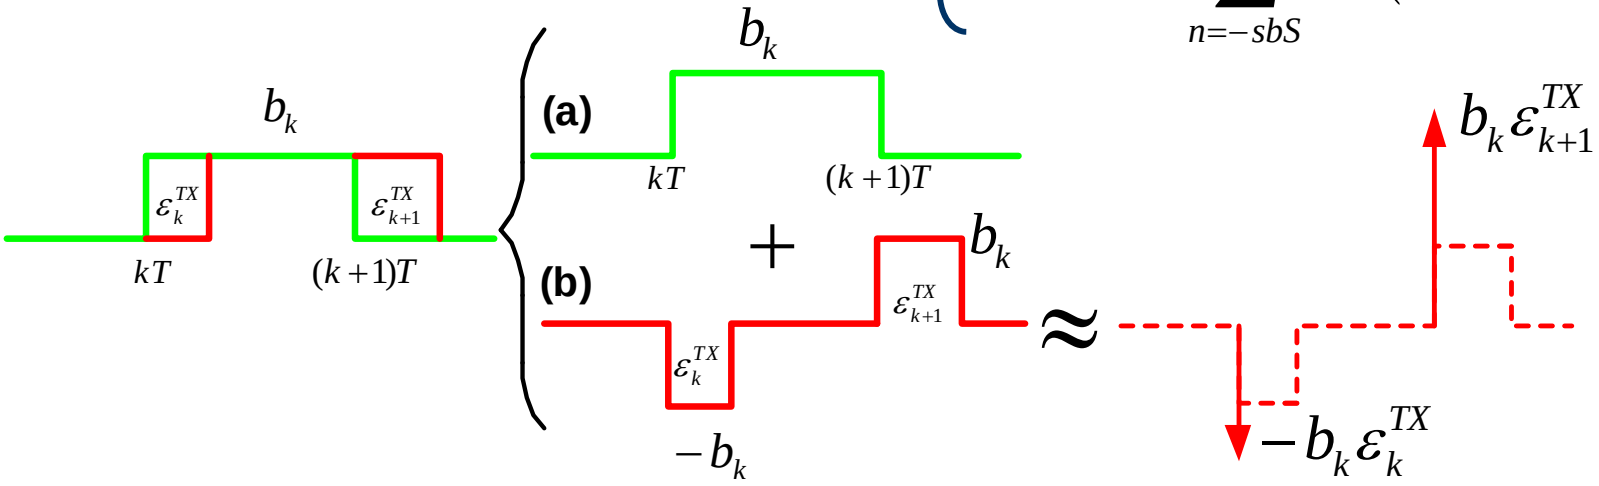

Deterministic Noise Enhancements

- Non-gaussian/uniform distribution exists
  - Circuit induced DJ has clustering effect
  - Duty cycle
- Incorporate into jitTX of the model

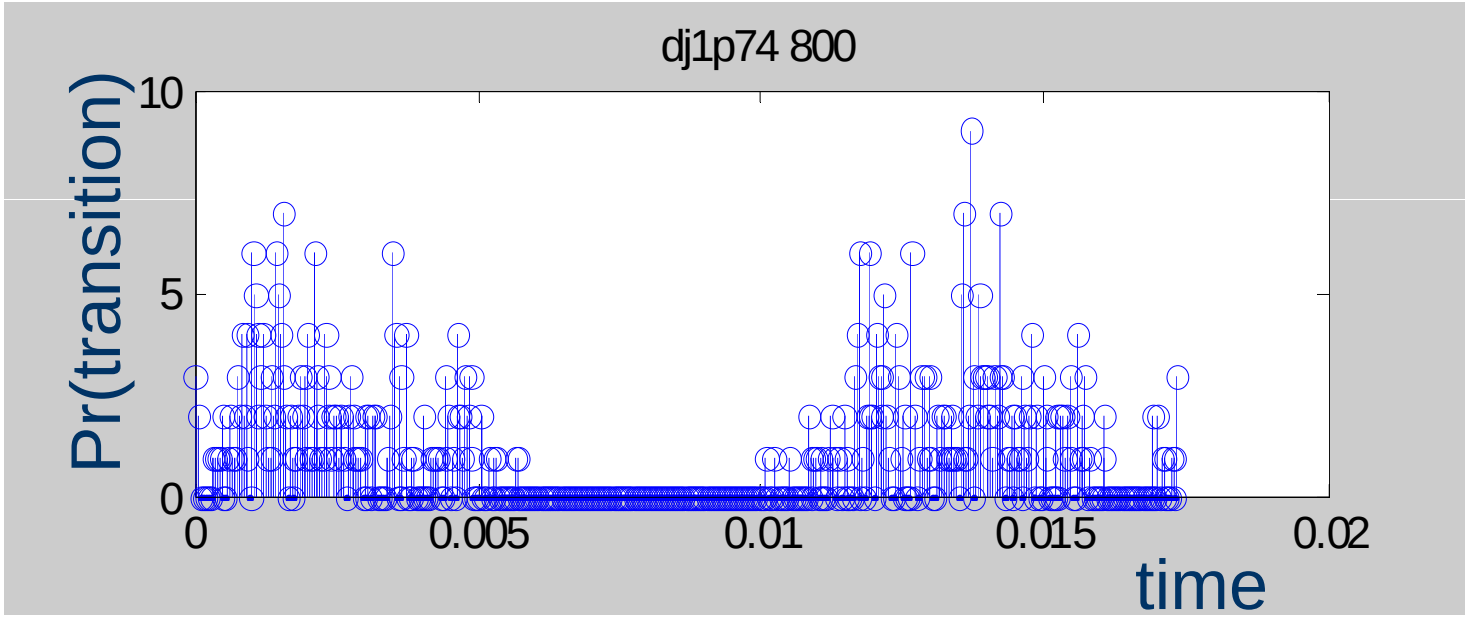

Signaling Power vs. Receive Margin

- 5mV receive margin for signaling power of 2.5mW (or 0.17pJ/bit)
  - 15Gbps through 8” channel + package

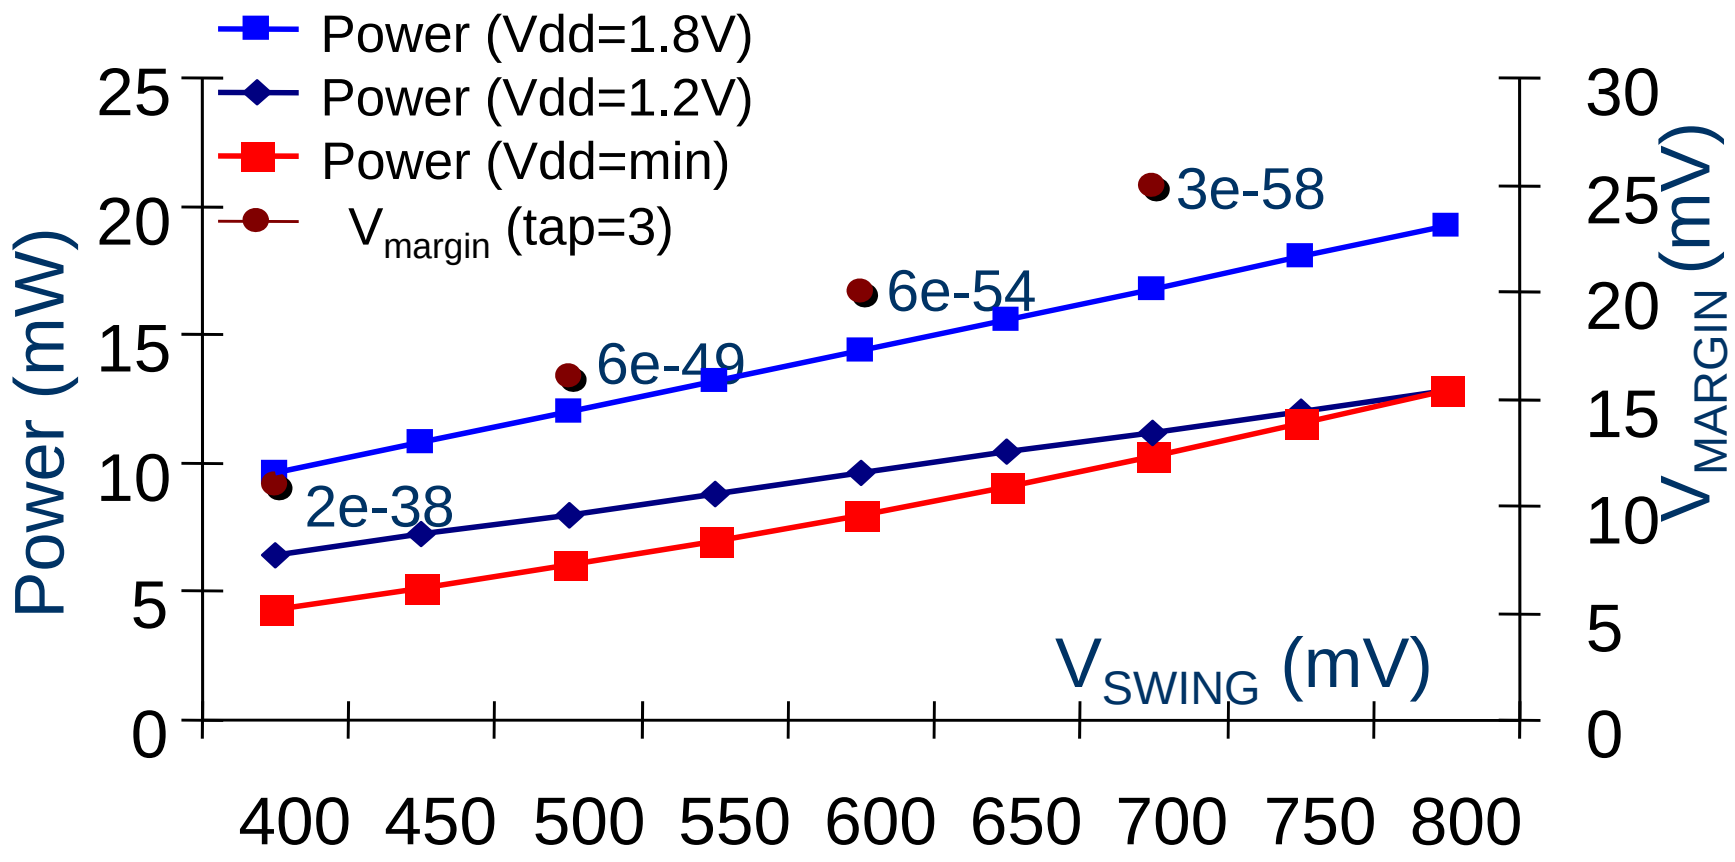

Outline

- Introduction to an I/O System
- Power-Centric Design Methodology
- Examples of Design Tradeoffs
- Conclusion

Sensitivity to Deterministic Jitter (DJ)

- Designers commonly minimize TX DJ.
  - <2% DJ leads to unnecessary energy cost
  - <1pJ/bit is possible
- Linear relationship between DJ and receive margin
  - 5mV margin per % DJ

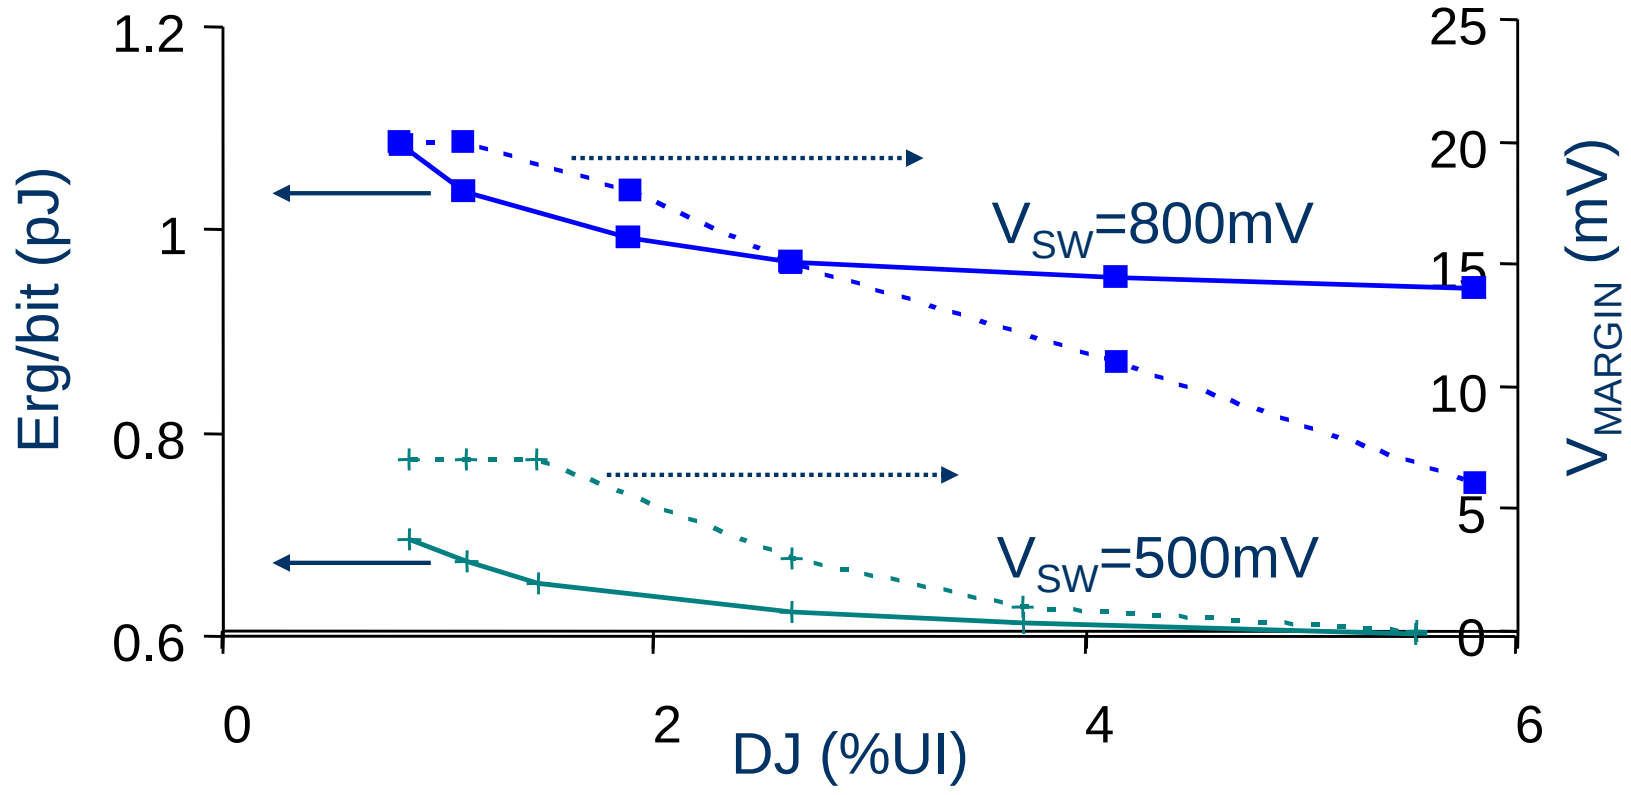

Tradeoff between Driver Architectures

- LCM driver dissipates more electronic power.
  - Increases more rapidly for low DJ.
  - <4% DJ leads to large energy cost.

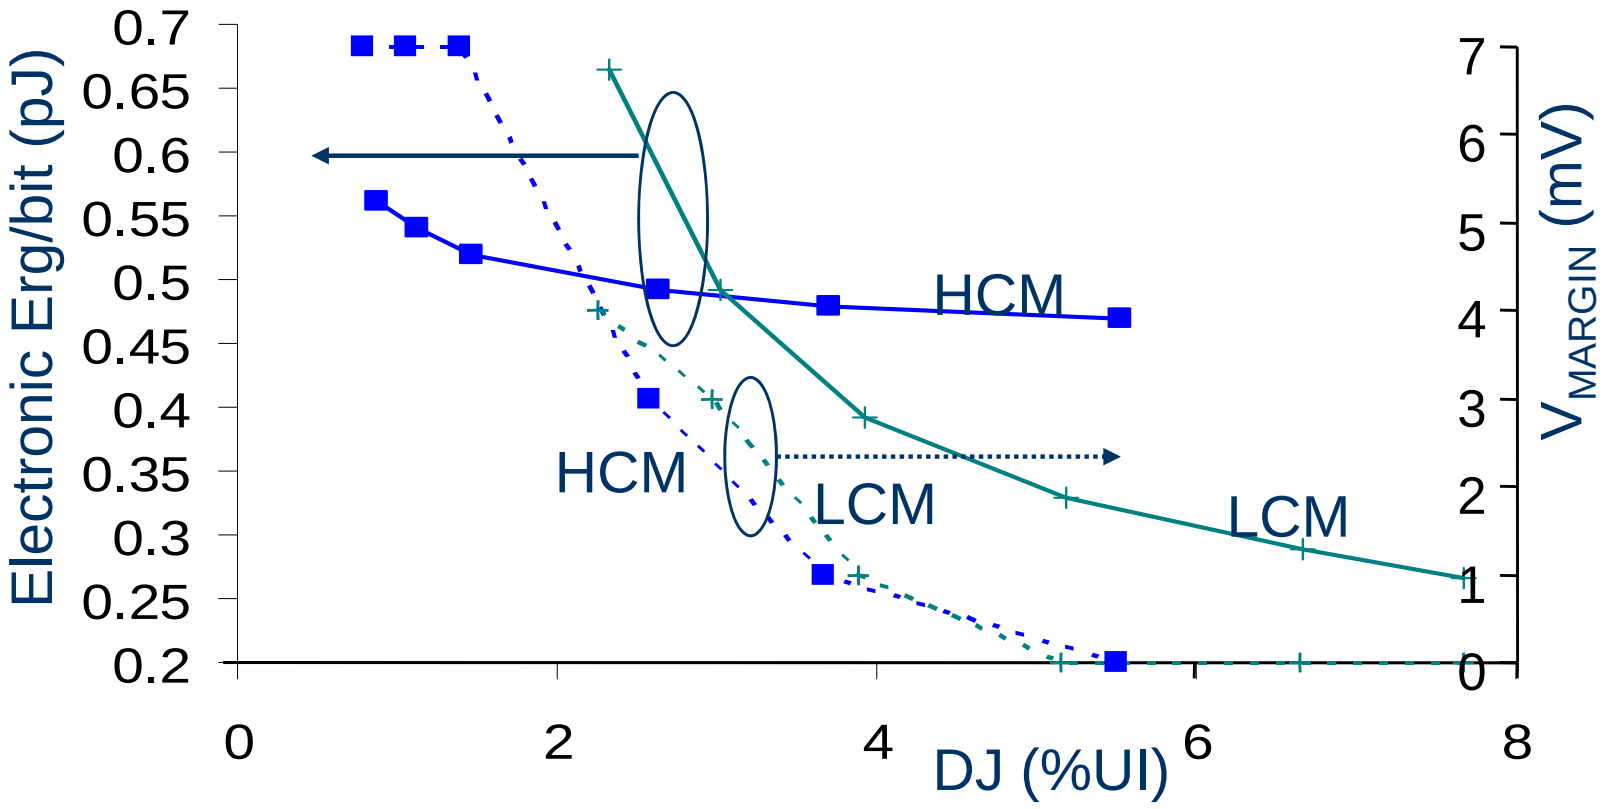

Multiplexer Position

- Mux earlier results in lower power (for the same DJ)
- Mux later results in higher max(f_{data})
  - When data rate is unconstrained

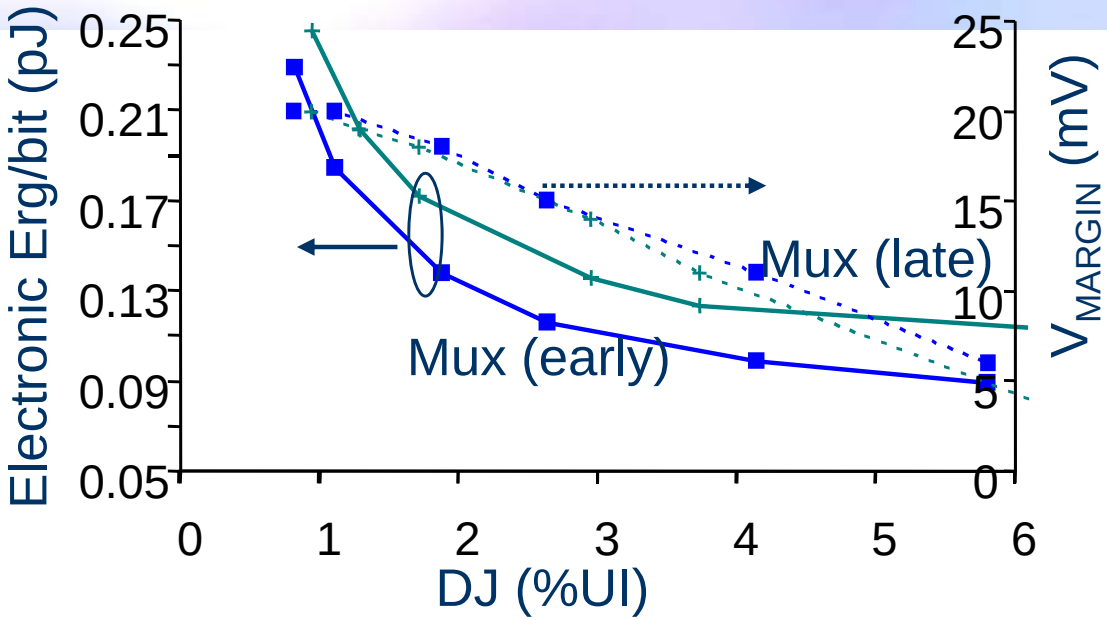

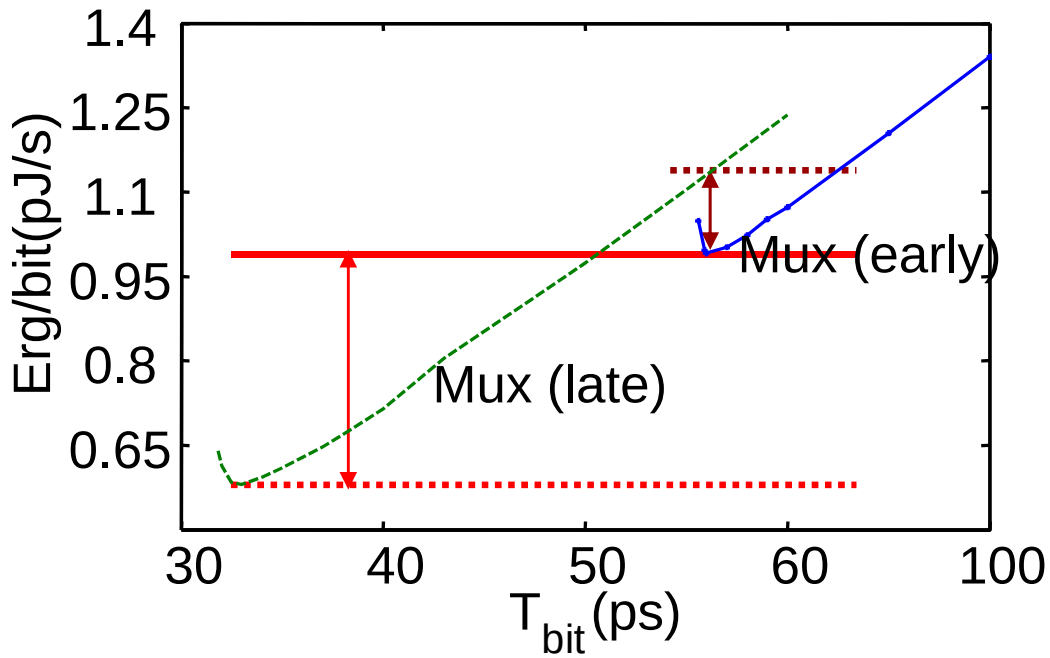

Energy Cost of Equalization

- Under-equalizing has large energy cost.
- Additional post taps do not have substantial impact.
  - Depends on the channel characteristics.
- Low energy cost (15fJ/bit-tap)
  - A better tradeoff than increasing signal swing (even the 3rd post tap)

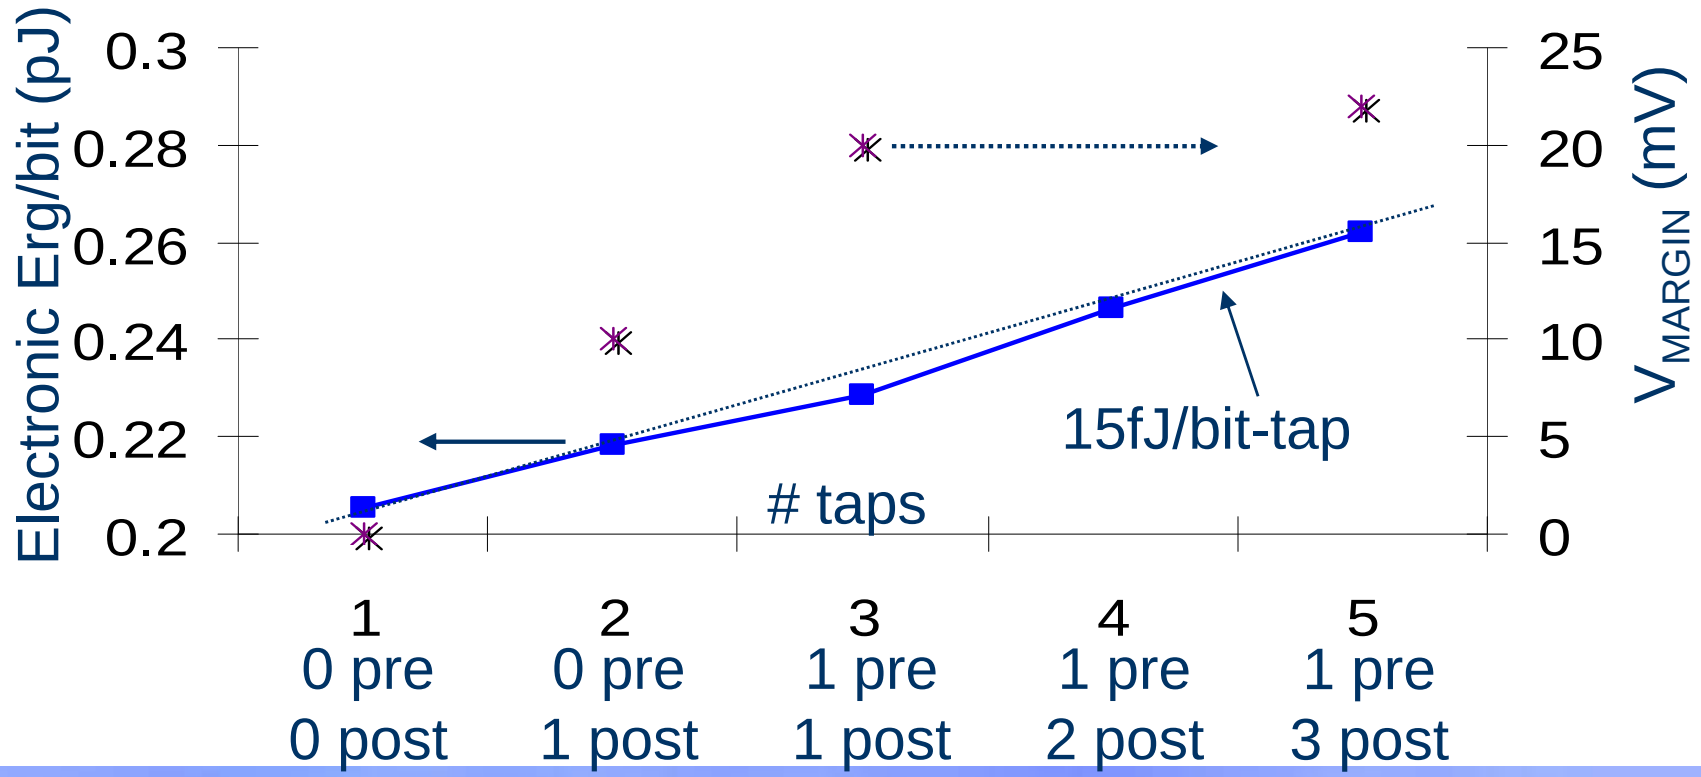

Conclusion

- Aggregate I/O bandwidth is scaling with logic density and speed
  - Energy/bit of I/Os must scale
- Large variation in I/O power
  - Over-design
  - Not being able to analyze the tradeoffs
- Need for optimization.
  - Optimize transistor-to-architecture
  - Incorporate a stochastic channel model
  - Handle mixed-signal circuits and I/O constraints