Multi-gigabit signaling with CMOS

William J. Dally - Massachusetts Institute of Technology

John Poulton - University of North Carolina @ Chapel Hill

Steve Tell - University of North Carolina @ Chapel Hill

This work supported under DARPA Contract No. DABT63-96-C-0039

Outline

- Intro -- The Bandwidth Gap
- Improved CMOS Signaling Systems
- Fundamental Limits & Possible Solutions
- Experimental 4Gb/s Signaling System
- Results and Conculsion

µProcessor I/O Bandwidth

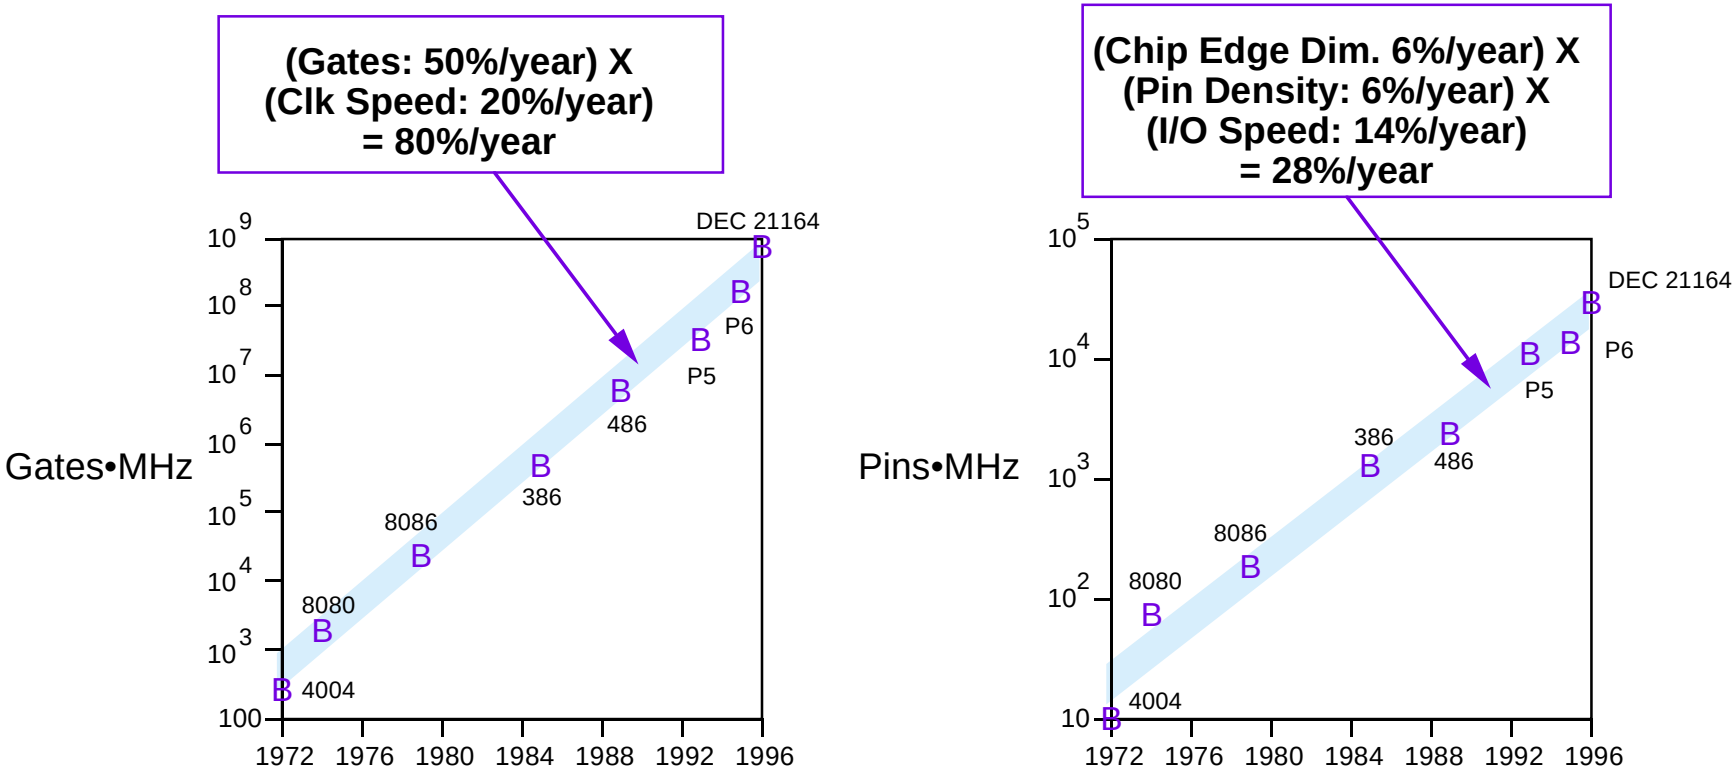

Connection between two graphs: **Rent's Rule**

$$\text{Pins} = K \bullet (\text{Gates})^{\alpha}$$

$$= 7 \bullet (\text{Gates})^{0.2}$$ for **Microprocessors**

Bandwidth Gap

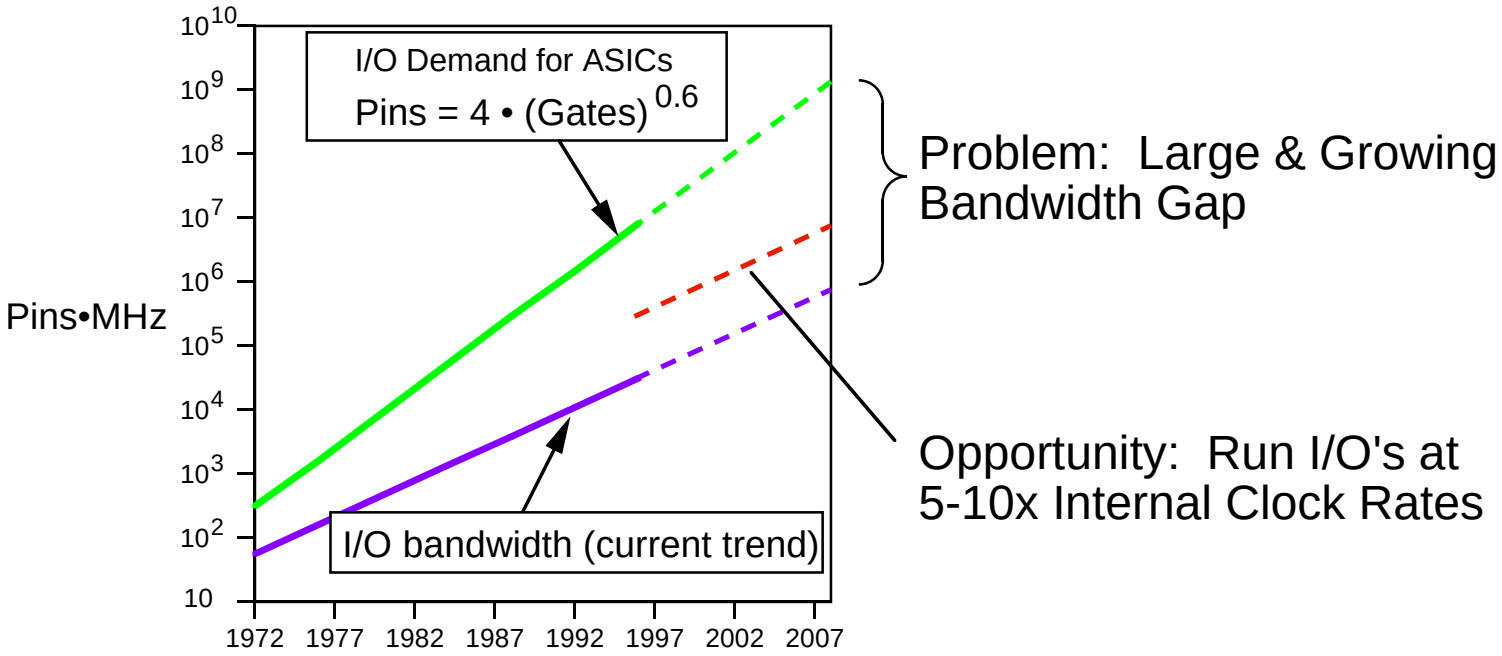

Bandwidth Gap

Current I/O signaling **anomolous** (I/O *slower* than internal clock)

Does it really take longer to send a bit from one chip to another than to do a 64-bit add??

Hostage to obsolete, poorly conceived signaling standards

Voltage-mode signals referenced to power supplies

Considerable recent progress:

Standards:

RAMBUS, IEEE1596/LVDS, IEEE1394, etc.

Special and Experimental systems:

Bull Serial Link (1992, 0.5µCMOS, 1Gb/sec)
Yang & Horowitz (1996 ISSCC, 0.8µCMOS, 2.5Gb/sec)
Fiedler, et.al. (1997 ISSCC, 0.5µCMOS, 1.06Gb/sec)
Chen & Baker (1997 ISSCC, 0.5µCMOS, 1.25Gb/sec)

Characteristics of 'Good' CMOS Signaling Systems

'motherhood'

controversial

- Differential

  Rejects external noise, most external sources common-mode
  Avoids reference-generation problem (bundled or local)

  Lemma: A differential signal pair can always be driven faster than two single-ended paths.

- Current-Mode

  Twin benefits:
  Decouples signals from power supply noise; avoids generating self-induced noise

- Controlled Edge Rates

  Edge time should never be shorter than half bit-cell time;
  Reduce reflections from impedance discontinuities

- On-Chip Termination

  A 'must' to avoid reflections from unterminated stubs

- Per-pin Timing (Receiver Clock) Recovery

  Avoids bundled clocking and difficult skew control problems

- Best Advantage from Available Devices

  CMOS NFETs good switches, accurate current sources
  PFETs poor switches, poor current sources, good resistors
  Argues for unipolar differential signaling

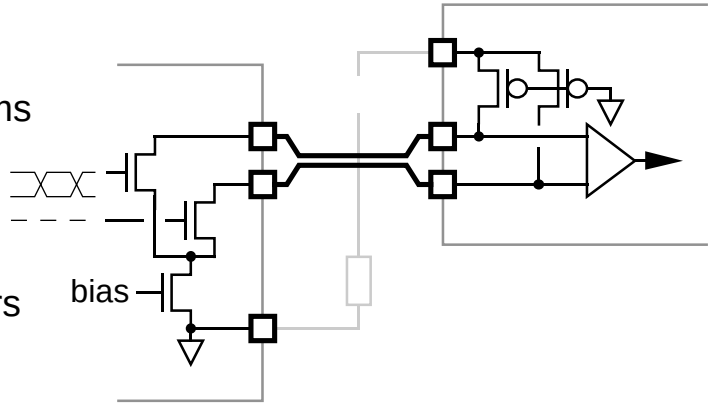

How Fast can CMOS Go?

- 0.5µ CMOS has FO4 inverter delay of about 180psec

  Can switch current onto a differential wire pair at 4Gb/sec fairly easily

- Recovering the data is harder!

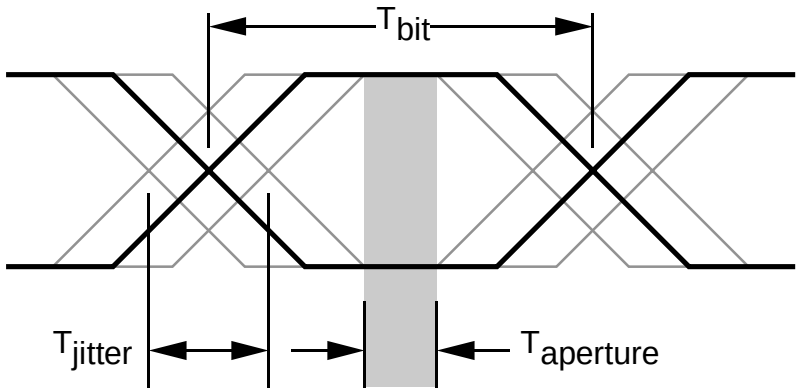

$$T_{bit} > T_{jitter} + T_{aperture}$$

- Critical circuit design issue:
  Generating precise, jitter-free multi-phase clocks

  10Gb/sec at 0.25µ??

Performance Limitations | Proposed Solutions

• Transmitter Speed	N-way multiplexed transmitter operates at 1/N bit rate
• Timing Uncertainty	Replica-biased differential timing elements (simulated jitter ~ 40psec in 0.5μ CMOS)
• Receiver Bandwidth	Multiplexed receivers Oversampling → Tracking (50psec aperture) Closed-loop timing recovery cancels skew & low-frequency jitter
• Reflections and Mismatches	Packaging design issue; on-chip terminations adjustable to a few %.
• Frequency-Dependent Attenuation (Skin Effect)	Transmitter equalization

Skin Effect

Frequency-dependent Resistance

K= 1.31X10-8 $\Omega s^{1/2}$ (round conductors)
K = 4.12X10-8 $\Omega s^{1/2}$ (thin strips)

$$R(f) = \frac{K f^{1/2}}{d} \quad (\Omega/M)$$

conductor dimension (width, radius)

Frequency-dependent Attenuation

$$A(f,x) = \exp\left(-\frac{R(f)}{Z_0} x\right)$$

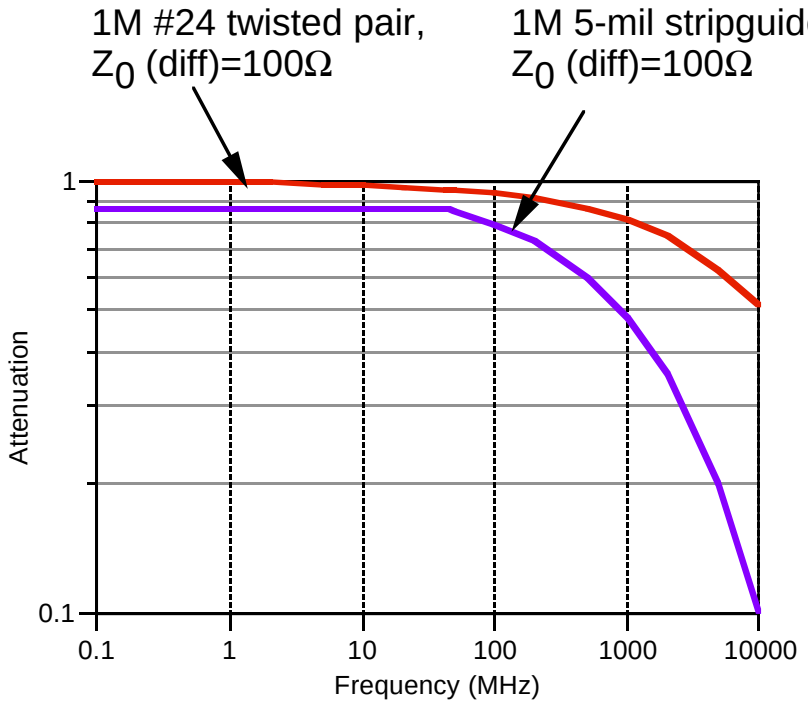

More attenuation from dielectric loss, radiation, package parasitics, lumped capacitance at load...

Intersymbol Interference

The 'Lone Pulse' Problem:

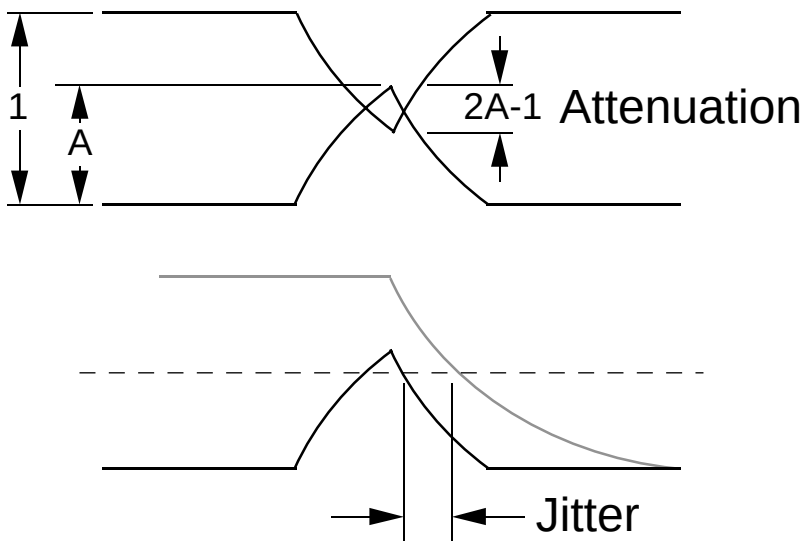

Equalization attenuates low frequencies at transmitter by A using (FIR) filter that complements line characteristics

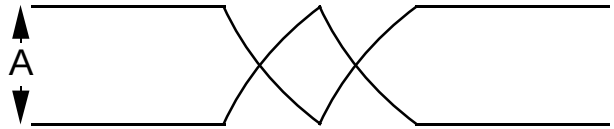

Transmitter Equalization

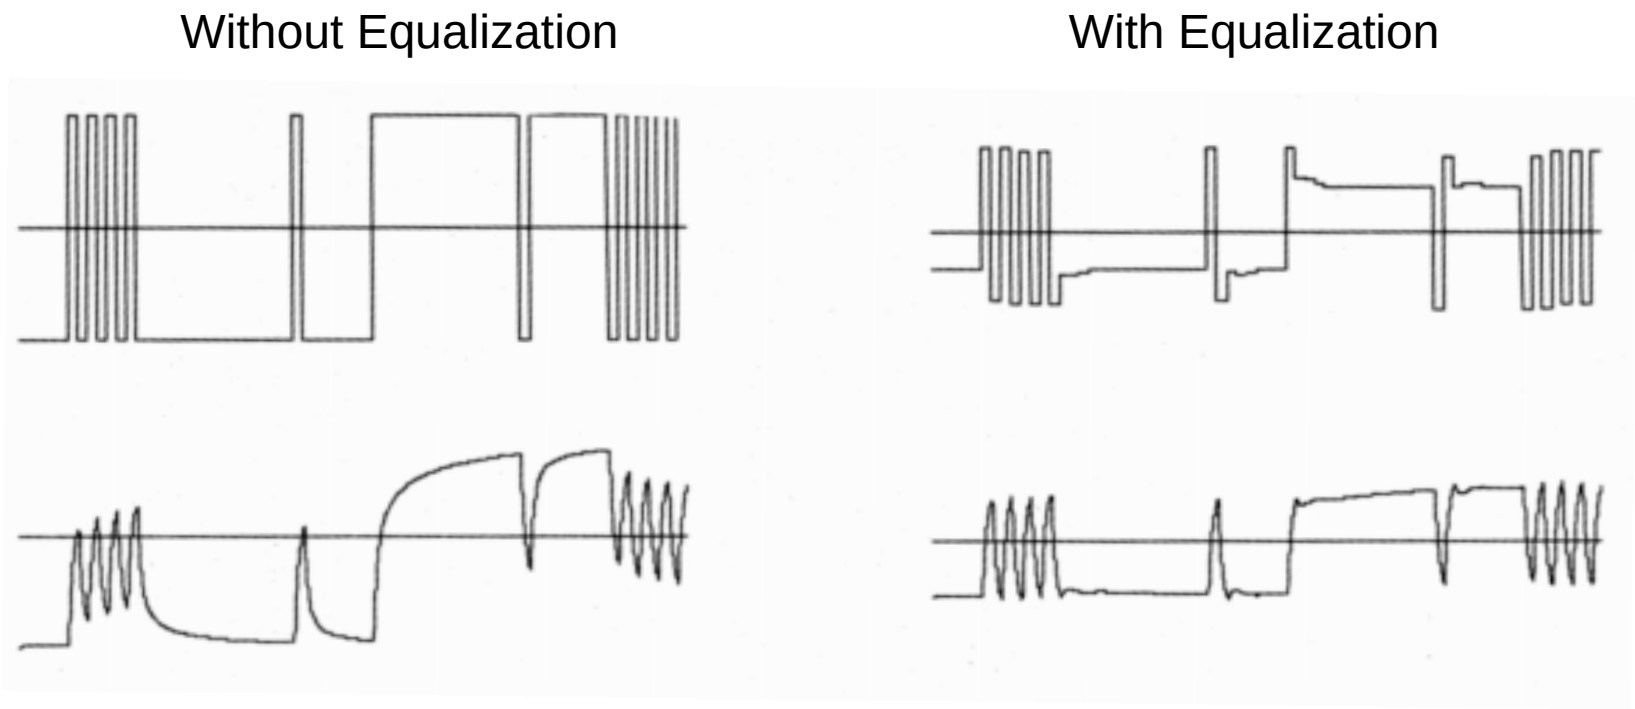

These are simulated results with a 5-tap FIR filter

Our implementation is a 5-tap *transition* filter:

Easier to implement and very nearly as effective

Multiplexed Transmitter

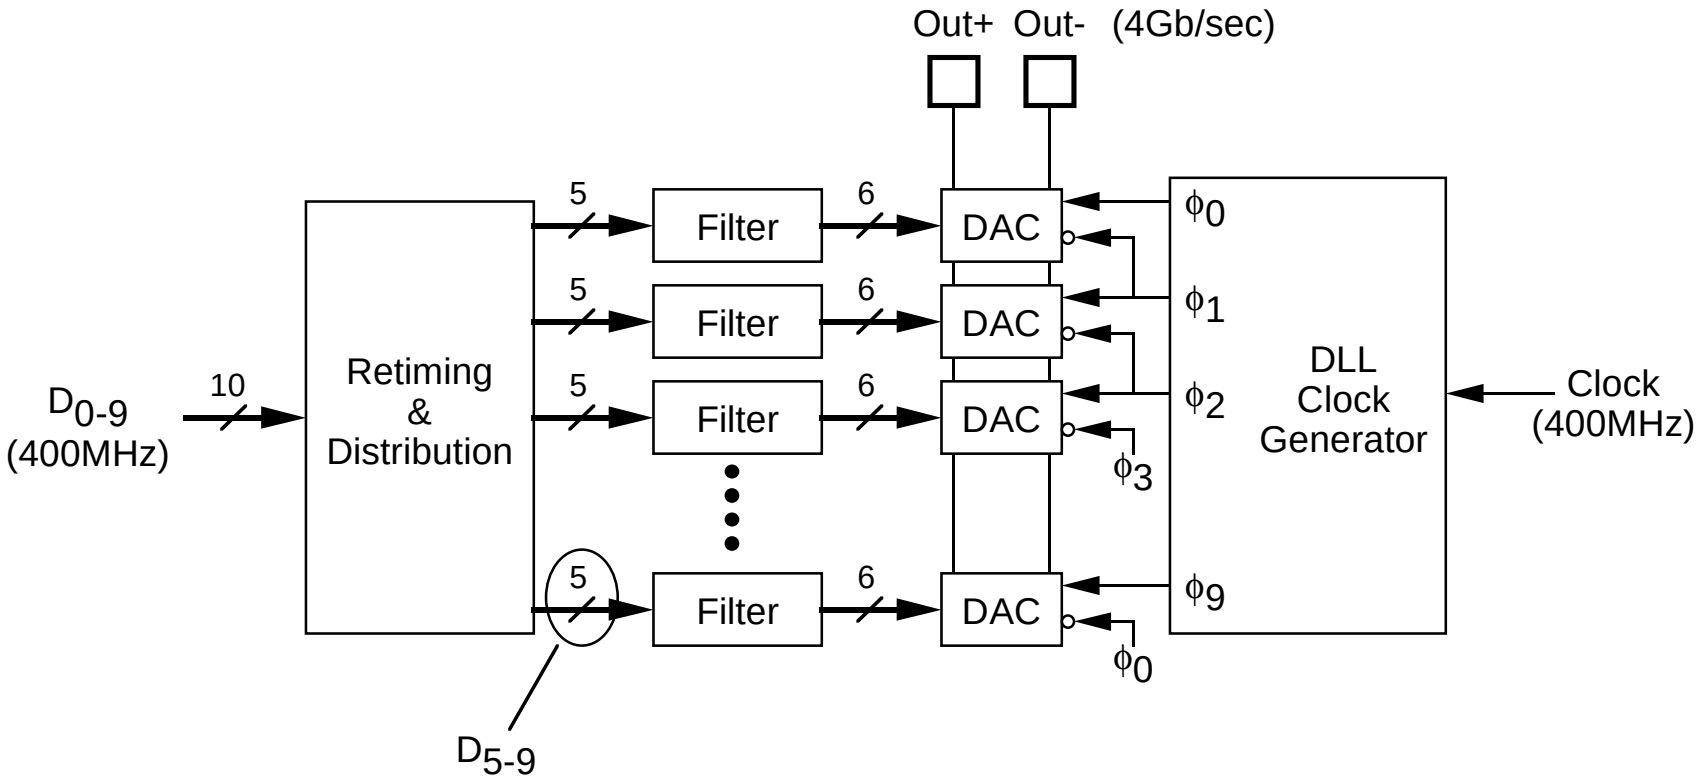

Transmitter Equalization Filter & DAC

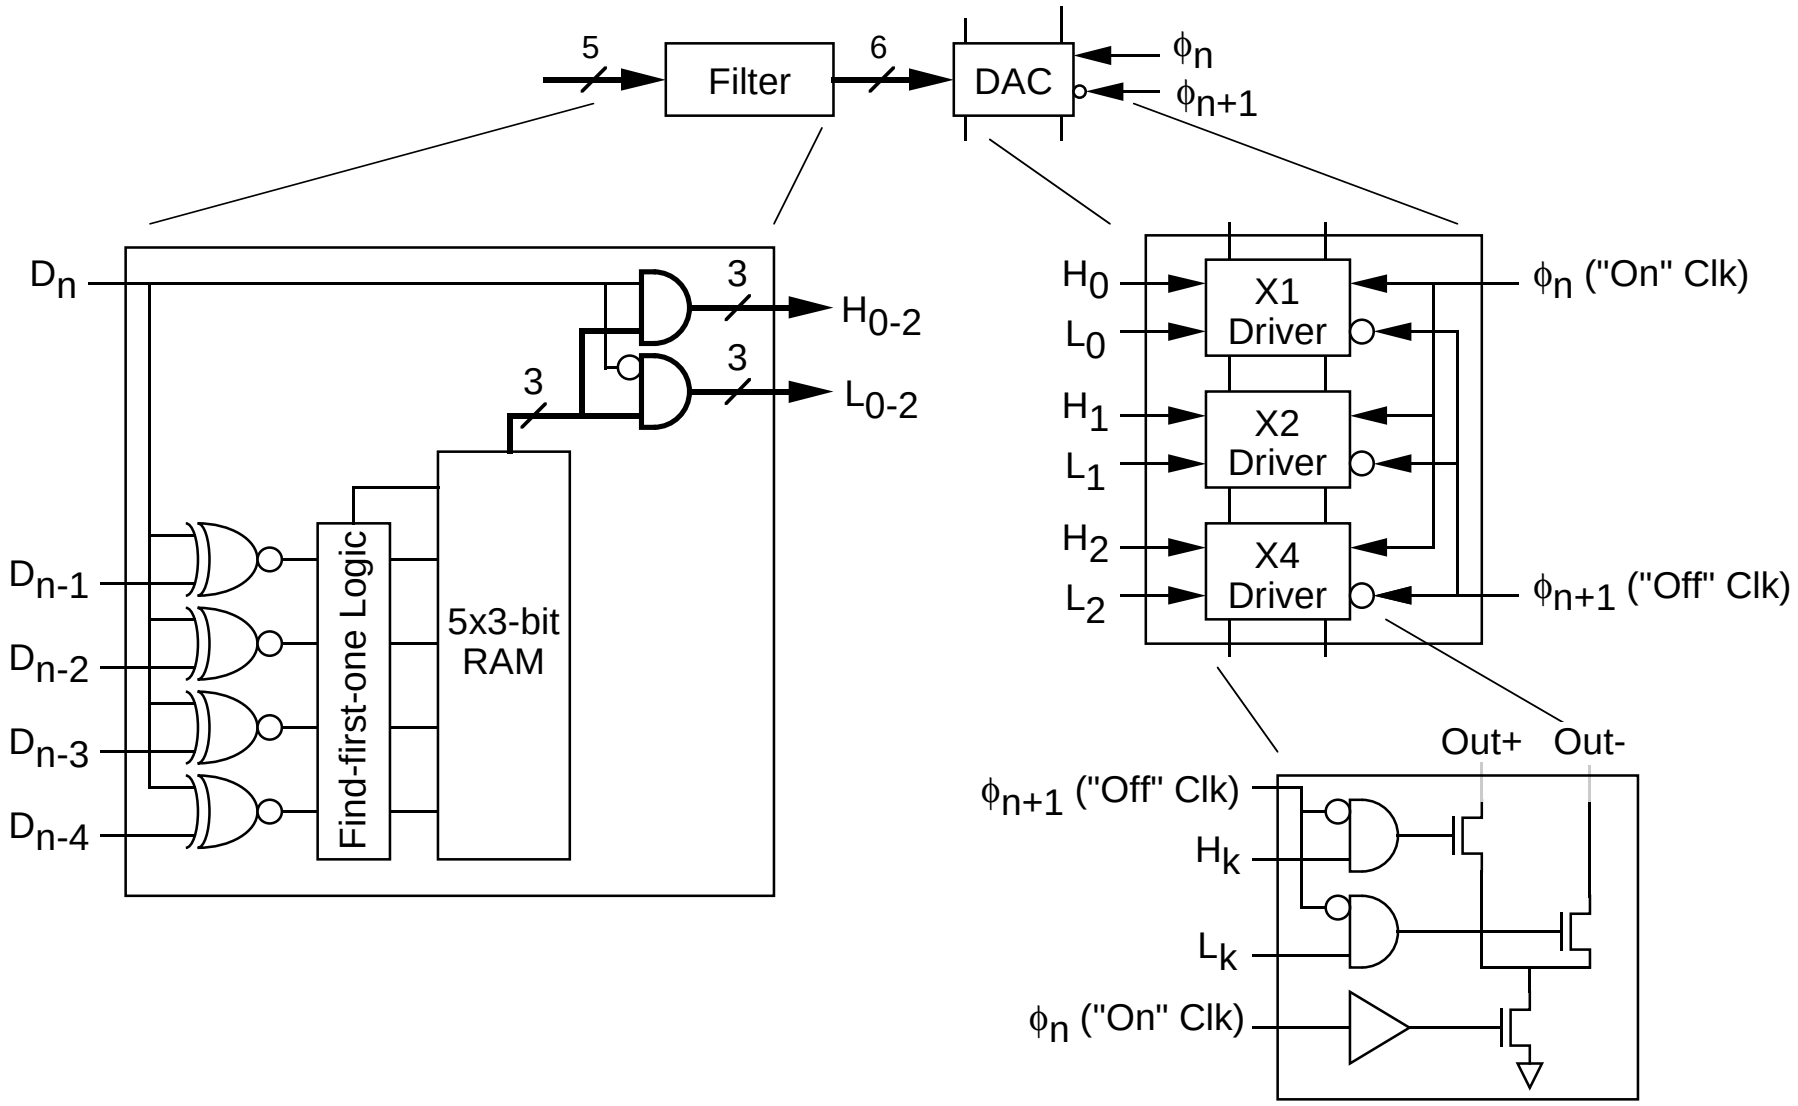

Receiver with Tracking Clock Recovery

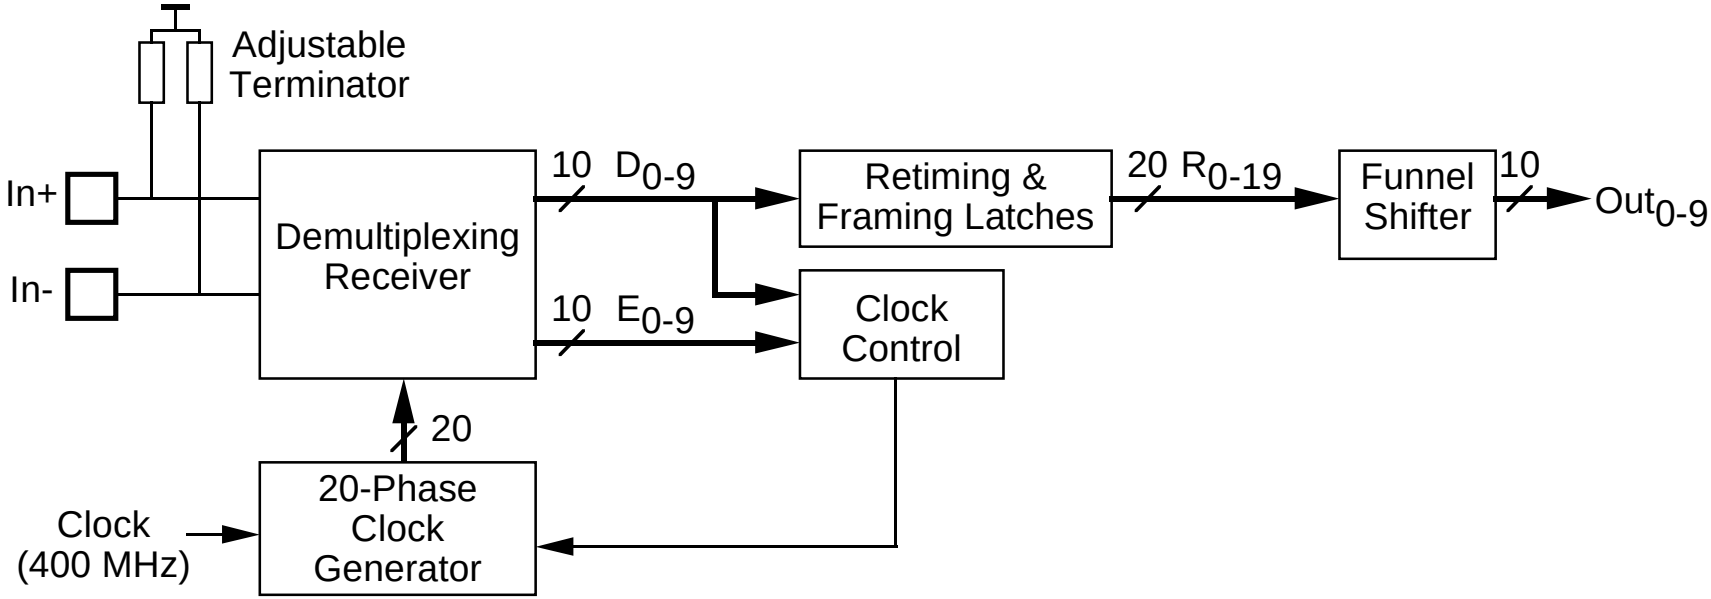

Demultiplexer

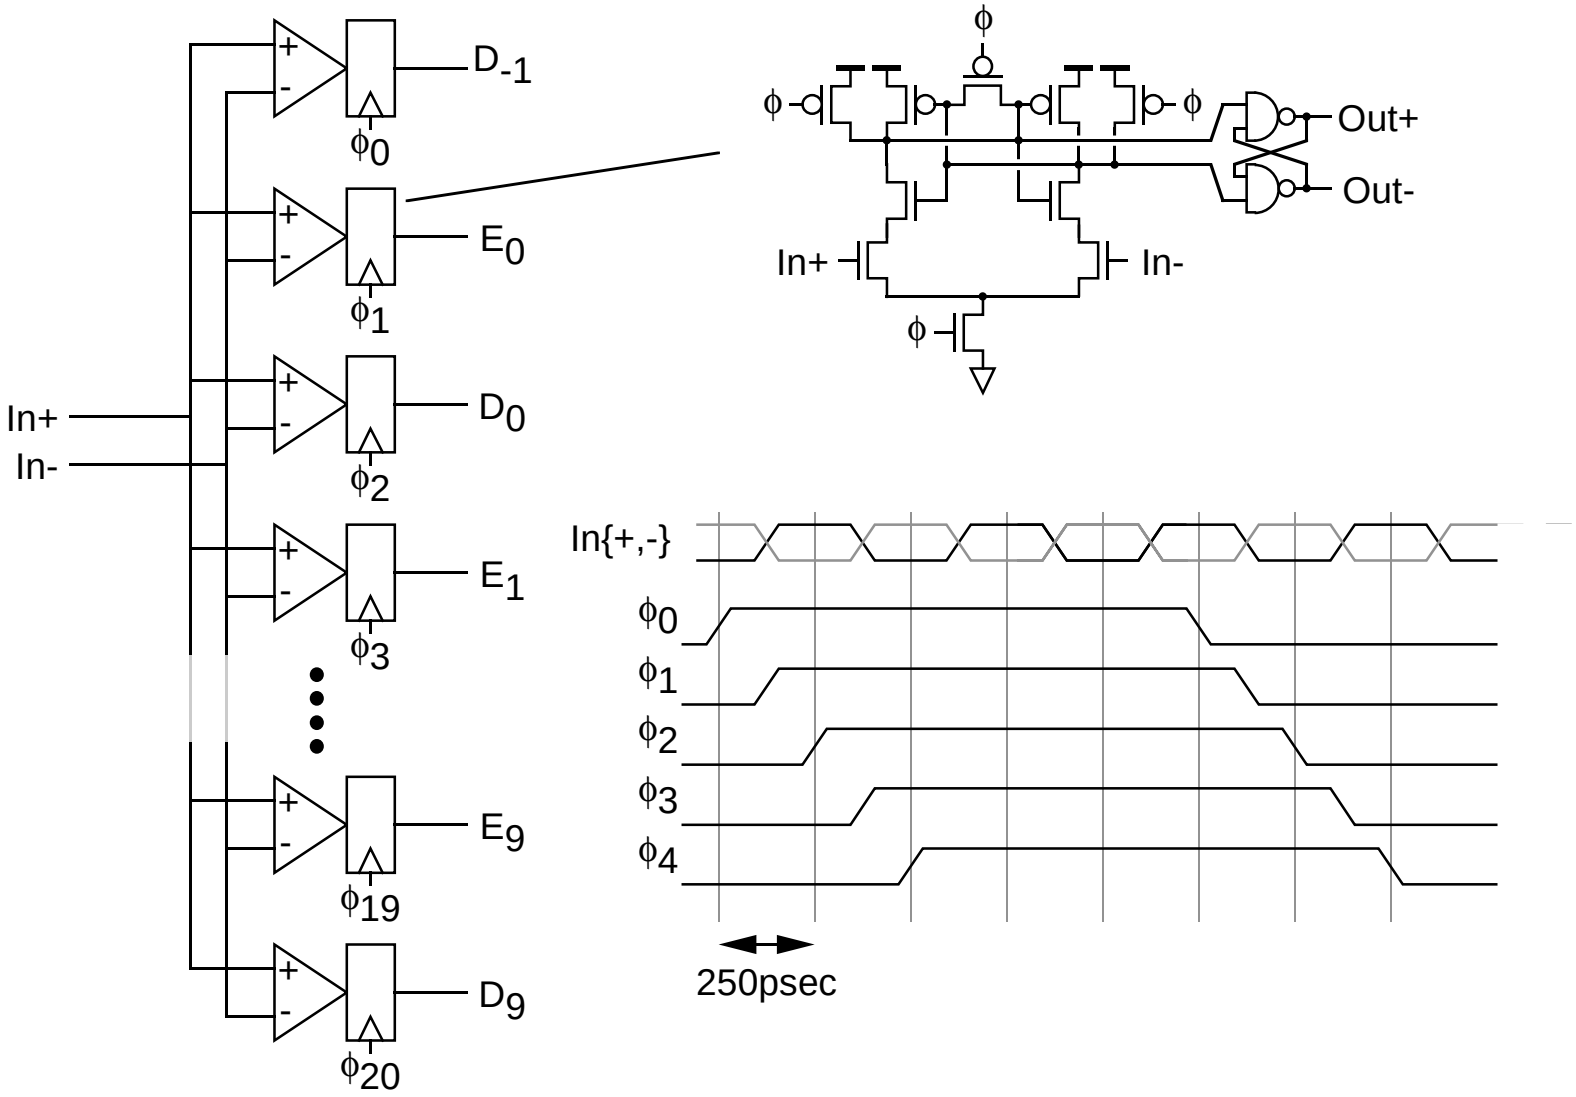

20-Phase Tracking Clock Generator

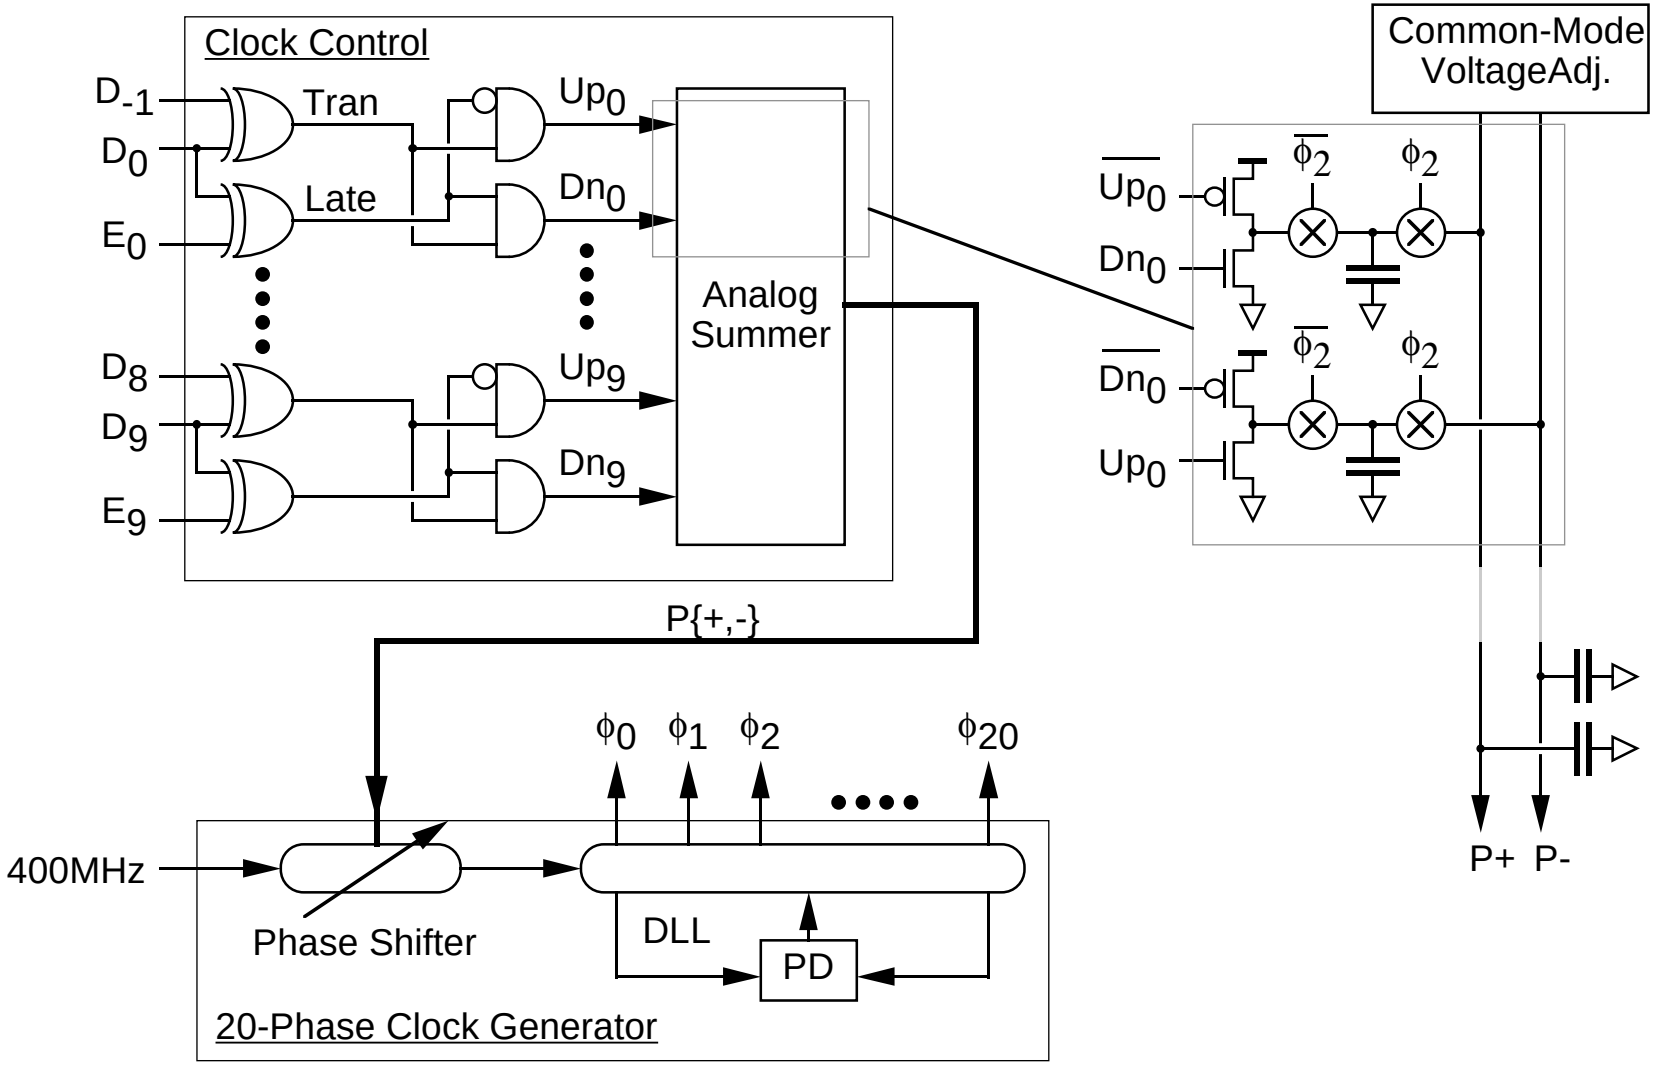

Clock Generator Details

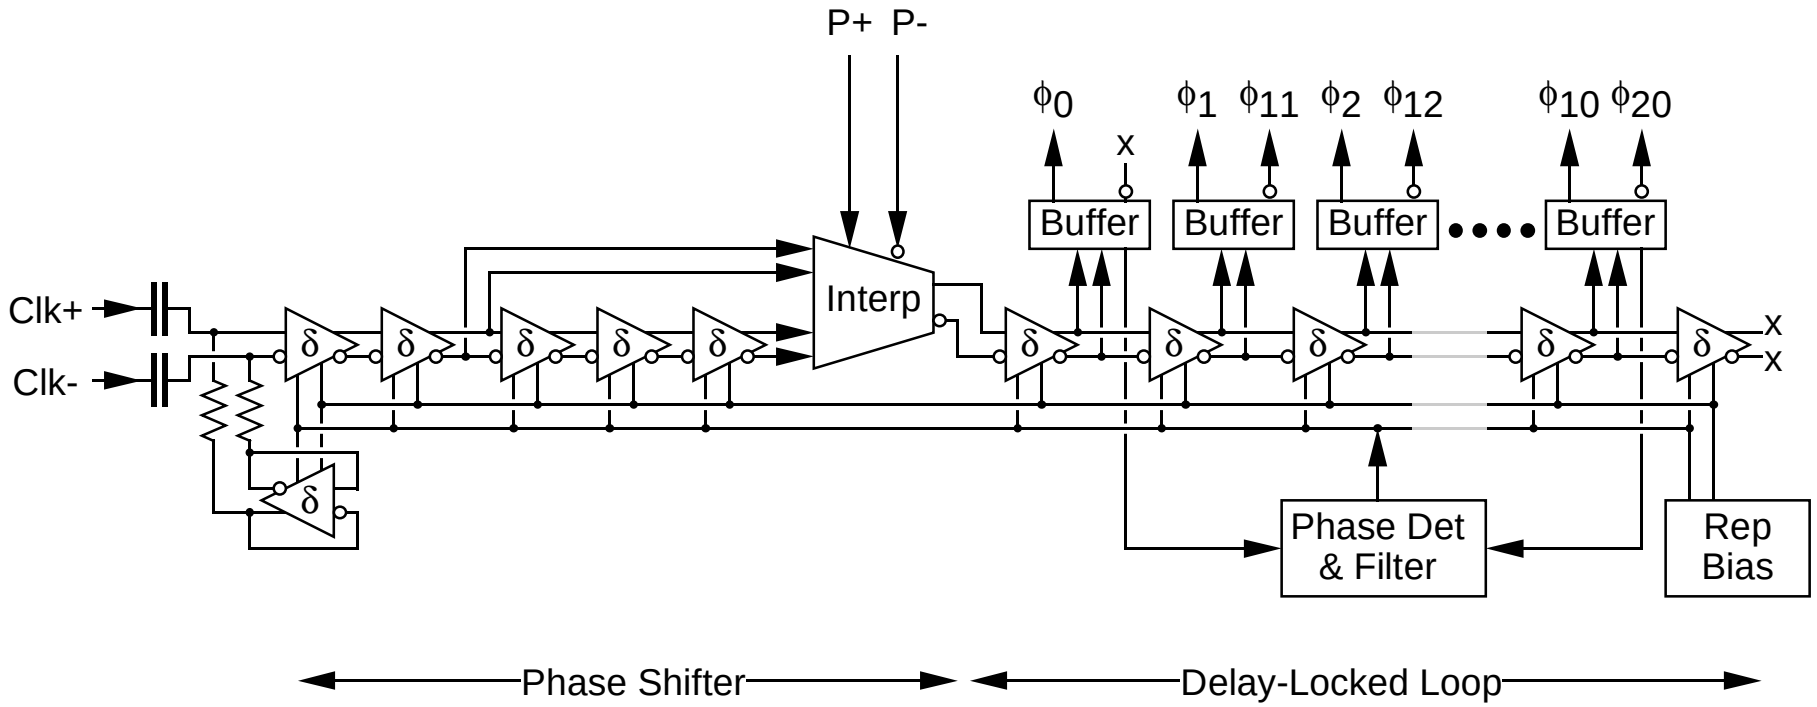

Test Chip #1

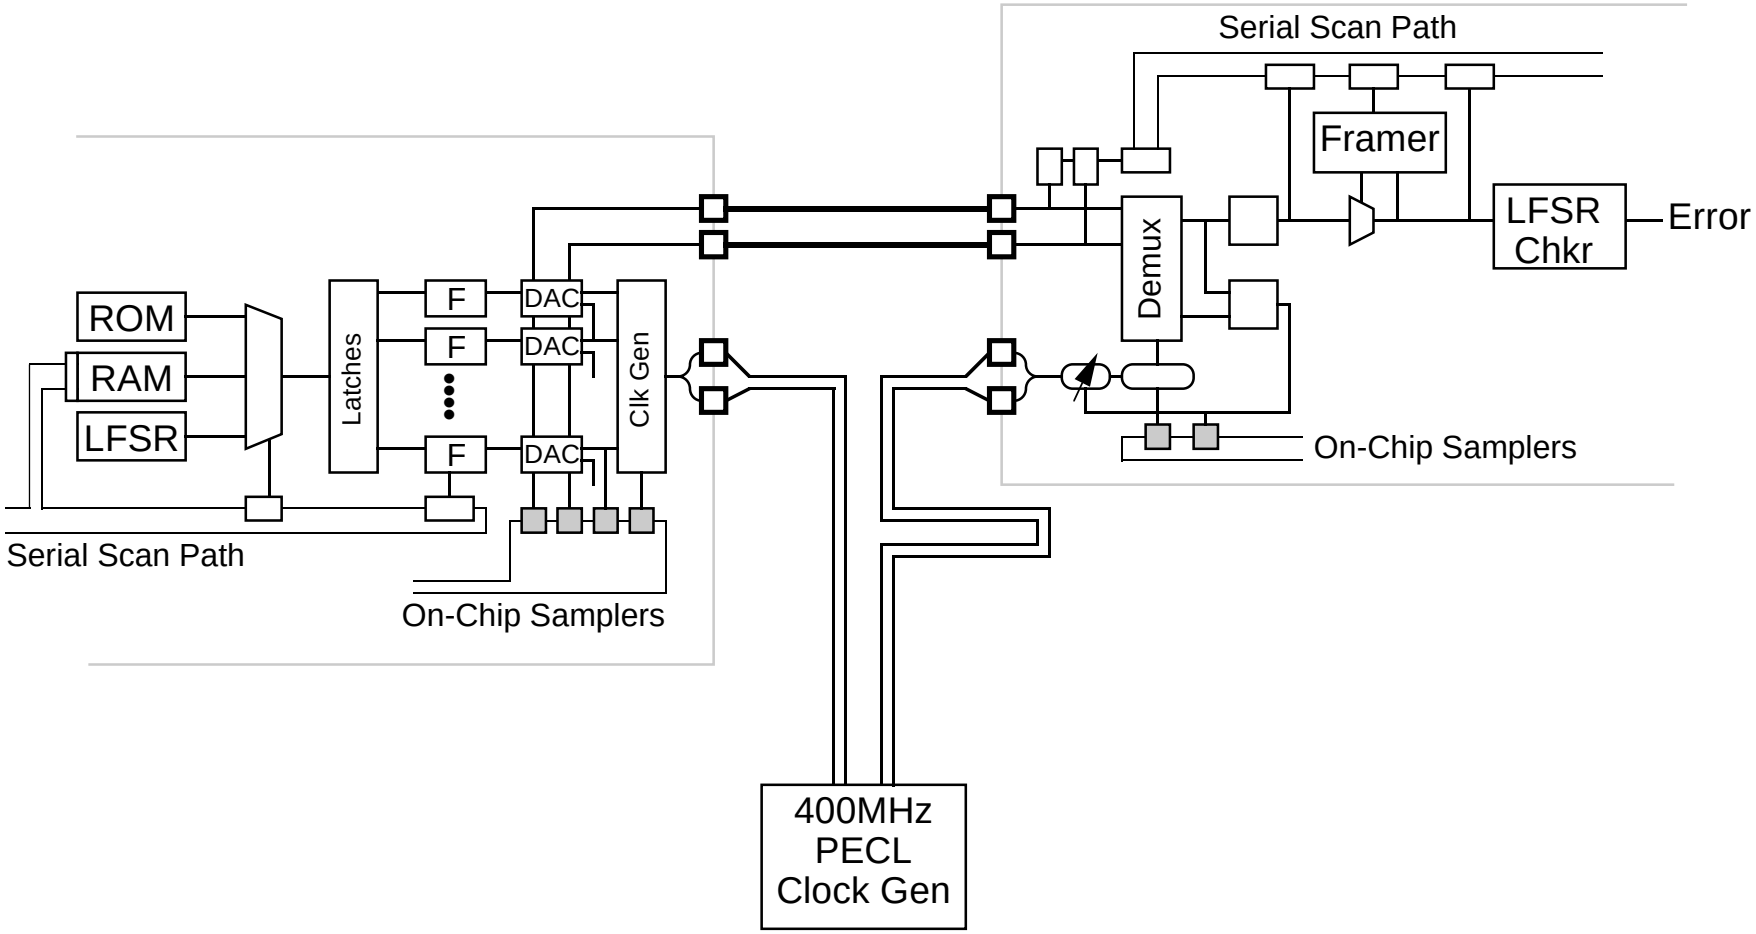
Serial Scan Path
Framer
LFSR
Chkr
Error
Demux
ROM
RAM
LFSR
Latches
F
F
F
DAC
DAC
DAC
Clk Gen
On-Chip Samplers
Serial Scan Path
On-Chip Samplers
400MHz
PECL
Clock Gen

Test Chip Waveforms

	Equalization Turned OFF	Equalization ON
Near		
Far		
Eye Diagram		

Conclusions

- Scaling of Gates (50%/year) much faster than I/O pins (12%/year)
- Can afford to spend 1000s of transistors for each off-chip signal
- Can & should clock off-chip signals faster than on-chip clock (not slower!)
- Today, relatively easy to signal at 1Gb/sec

  Differential
  Current-mode
  On-chip termination
  Per-signal clock recovery

- Use all those transistors to attack fundamental problems

  Inter-symbol interference due to skin-effect frequency-dependent attenuation

  Echo cancellation?
  Simultaneous bi-directional signaling?
  > 1 bit / clock?

Future Work

New test chip Fall 1997

- N-way muxing transmitters with lower N
- Simpler transition filters (5 elements probably more than needed)
- Compare oversampling with tracking clock recovery
- Alternative timing circuits
- Coding and error correction strategies
- Exportable 1Gb/sec version in 0.5μ CMOS

In what year, and why, will optical interconnect become the dominant technology in high performance digital systems for lengths > 10cm?

Interconnect Length > 10M : Already the dominant technolgy

Interconnect Length 1-10M : Electrical and Optical Signaling Both Contenders

	Launch Cost	Cable Cost	Isolation	Power/Area
Optical	High	Low	Easy	High
Electrical	Low	High	Difficult	Low

Current work rapidly reducing optical signal launch costs, so may become the dominant inter-cabinet signaling technology soon

Distances < 1M??

Today's 2nd-level packaging paradigm: board delivers power, provides signal interconnect; heat removed above board

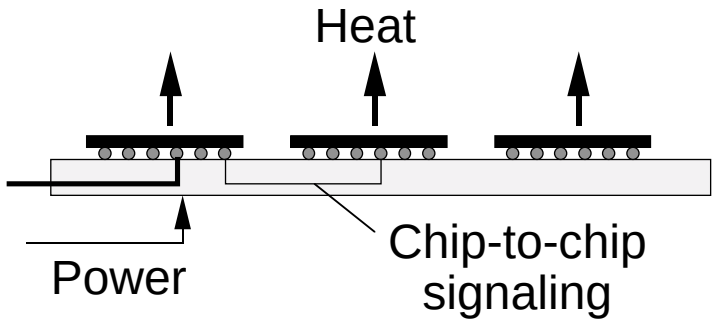

Signal scaling properties:

	1997	Scaling	2007	
Signal pins/chip	200 250 (pairs)	1.12/year	600 750 (pairs)	← Conventional ← Aggressive
Signaling speed	100Mb/sec 1Gb/sec	1.14/year	400Mb/sec 4Gb/sec	
Power per signal	50mW	0.8/year	5mW	
Signal width/pitch	125/250µ 40/100µ	0.94/year	60/120µ 20/50µ	
Length for Attn=0.5 (Aggressive)	~ 1M		~ 0.2M	

Distances < 1M??

To replace electrical signaling with optical will require:

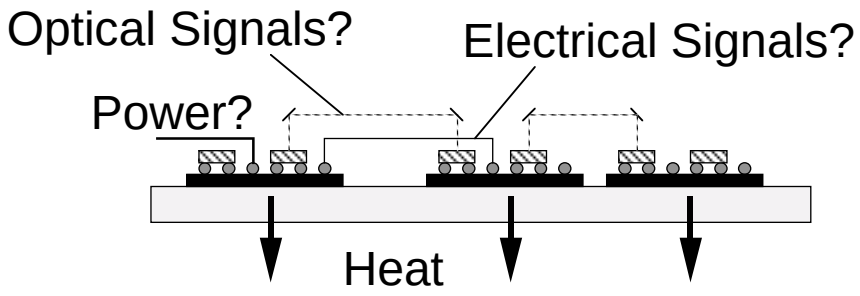

Graceful way to deliver electrical power and (legacy) electrical signals to the top of composite chips (combine with SEEDs processing?) and get heat out the bottom.

Build entirely new design-flow infrastructure, replacing well-established infrastructure for electrical signaling:

- Mixed-technology design
- Simulation
- Test

Develop some way to emulate PC-board wiring in the optical domain (this may be an opportunity--processing in the optical domain??)

Distances < 1M??

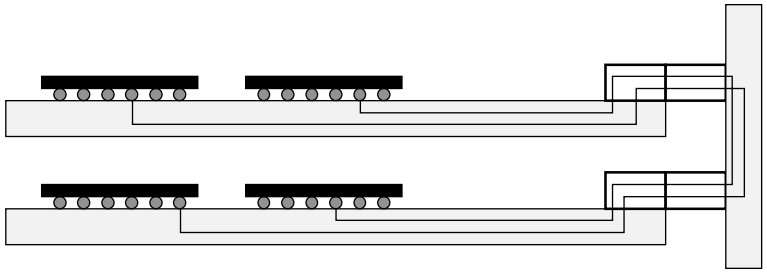

Optical 'Backplanes'

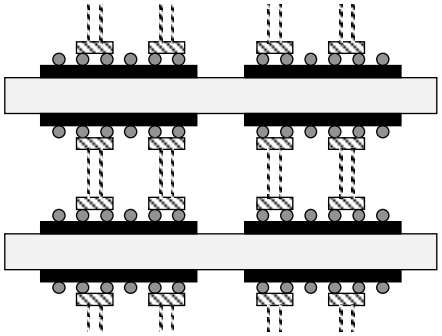

Daughter-card/Backplane paradigm provides interconnection density that scales only linearly

Connector pin density scaling even more slowly than board signals

Long routs to edge of card, then back; impedance discontinuities; expense; reliability; etc., etc. (litany of connector woes)

Optical interconnect provides another dimension to rout in

Energetically very favorable; potentially very high signal density

Don't have to solve all of the optical interconnect problems at once!

Could become an important technology within 10 years??

Recent Progress

Transmitter Pre-Emphasis (@ 1.06Gb/sec) - LSI Logic

Fiedler, A., R. Mactaggart, J. Welch, and S. Krishnan, "A 1.0625Gbps Transceiver with 2x-Oversampling and Trasnmit Signal Pre-Emphasis", ISSCC97, pp 238-239.

4Gbps Signalling - Yang & Horowitz, Stanford Univ

Private correspondence

Tracking Clock Recover - Symbios Logic, Inc.

Chen, D-L, and M. Baker, "A 1.25Gb/s 460mW CMOS Transceiver for Serial Data Communications," ISSCC97, pp 242-243.