NC STATE UNIVERSITY

Bandwidth, Bandwidth, Bandwidth

Paul Franzon,
Department of Electrical and Computer Engineering
NC State University
919.515.7351, paulf@ncsu.edu

Outline

- Application and the demand for bandwidth
- The nature of wires
- Ways to improve wires
- 3D Architectures
- Ways to reduce bandwidth

Microprocessors

- Multicore + other innovations ➔ 1 TFLOP uProcosser with 1 TB/s of memory bandwidth by 2010

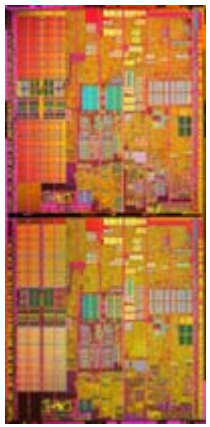

Intel

MULTICORE AND REVERSE SCALING

	2004 Baseline	Multi-core Approach	Reverse scaling	Reverse scaling
Frequency	4 GHz	8 GHz	8 GHz	4GHz
No. of Cores	1 Core	4 Cores	16 Cores	16 Cores
Core rel. IPC	1	1	0.5	1
Total Flops	32 GFlops	256 GFlops	512 GFlops	512 GFlops
Supply	1.2V	1.0V	1.0V	1.0V
Power	84W	233W	233W	117-163W
Bandwidth requirement	32GB/s	256GB/s	512GB/s	512GB/s

Future microprocessors and off-chip SOP interconnect
Hofstee, H.P.;
Advanced Packaging, IEEE Transactions on [see also Components, Packaging and Manufacturing Technology, Part B: Advanced Packaging, IEEE Transactions on]
Volume 27, Issue 2, May 2004 Page(s):301 - 303

Graphics Engines

- **Should hit 1 TBps before 2010:**

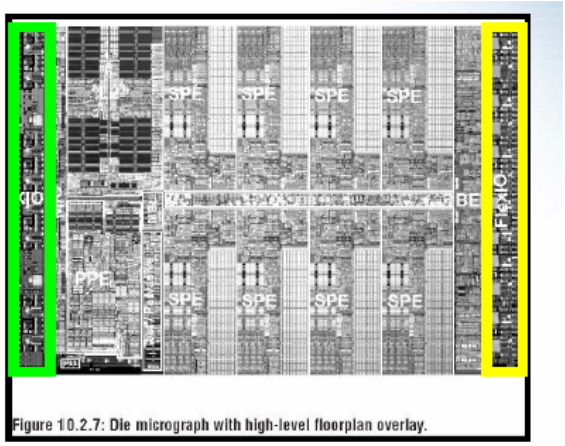

Cell CPU with 25.6GB/s XDR Memory Interface

c/- Rambus

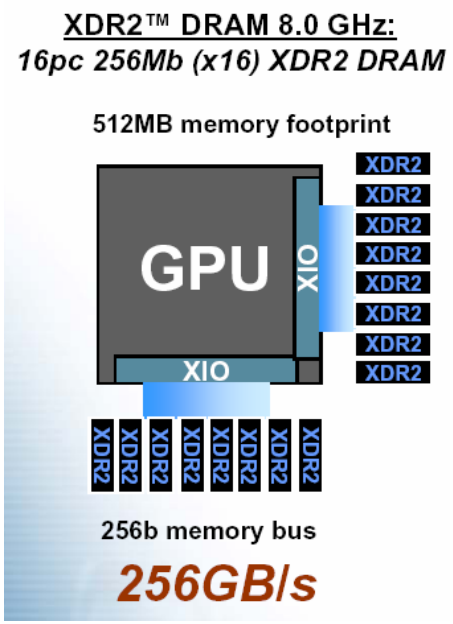

100M+ Triangles/sec
10M+ Triangles/image
1-10 TBps

2006 2008? 2010

Network Processors

- Scaling of 7-layer processors

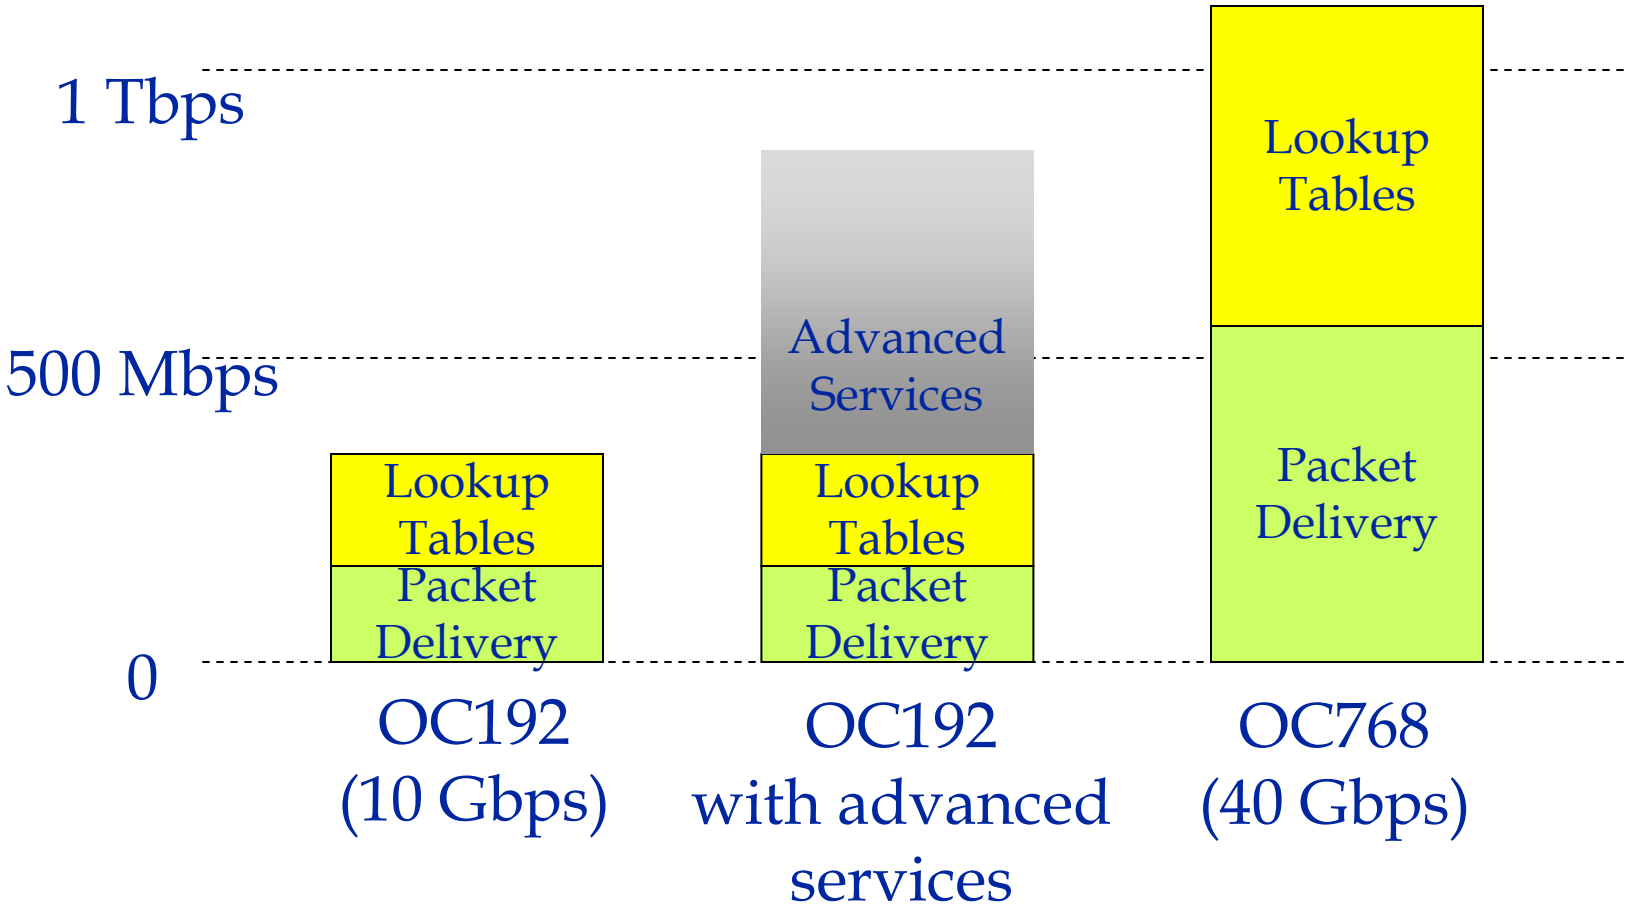

Producing 1 TBps of bandwidth

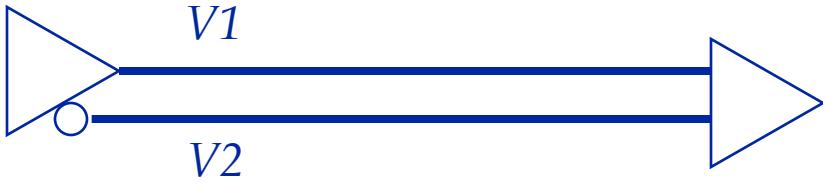

Using differential pairs

- **400 pairs at 20 Gbps**
  - 800 pins
  - + 400 pins for “control”
  - + 800 pins for power/ground
  - ➔ 2,000 pin package with 800 “differential” pins
  - @ 100 mW / pair ➔ 40 W for I/O
  - @ 30 mW / pair ➔ 12 W for I/O
- **At limit of physical achievability for current packaging technologies**
- **Beyond current technologies to control noise and jitter**

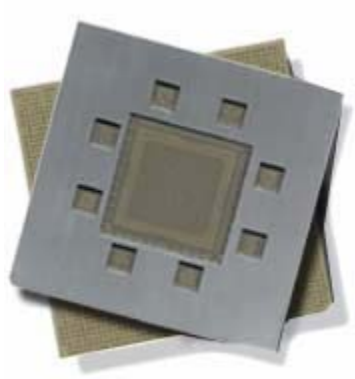

c/- EIT

Making Wires

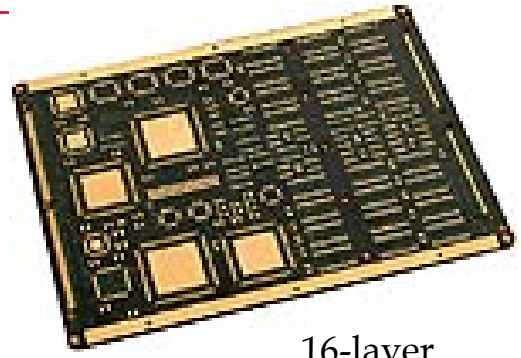

16-layer
150 um wire
PCB

- **PCB Cross-section:**

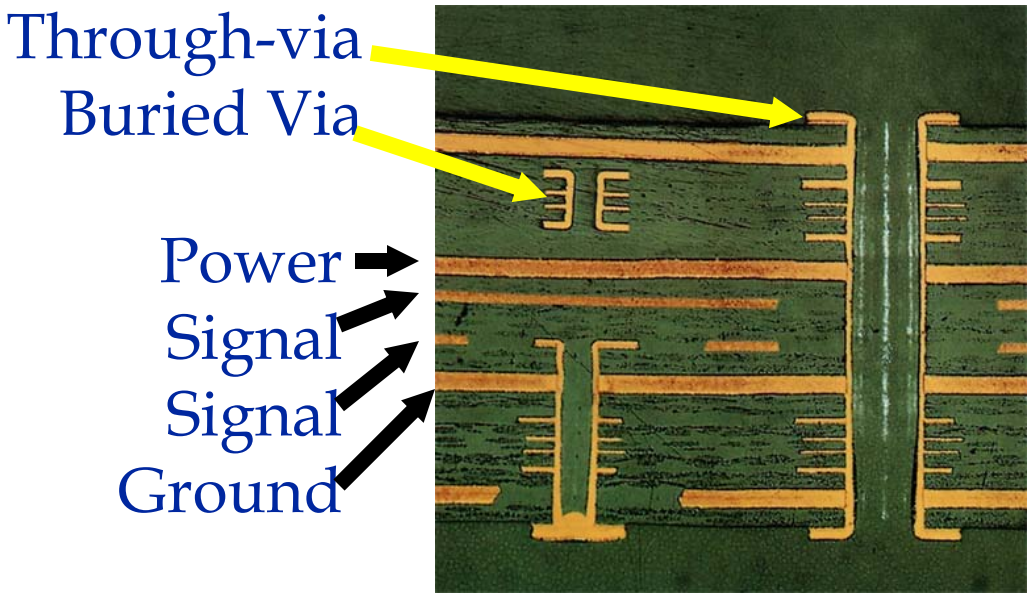

Traces:
15 – 30 um thick Cu

FR4 dieletric

Problems with current interconnect technology

- **Vertical registration difficult below ~25 um**
  - ➔ Limit on wire width of 25 – 75 um
  - ➔ Limit on via opening to about 25 um
- **Thickness limited to about 0.1 inches**
  - ➔ About 20 layers
- **High-frequency losses are high**
  - Skin Effect
  - Dieletric loss
- **Noise sources are mounting**
  - Crosstalk
  - Reflection noise from via “stubs”
- **And, connectors throttle bandwidth**

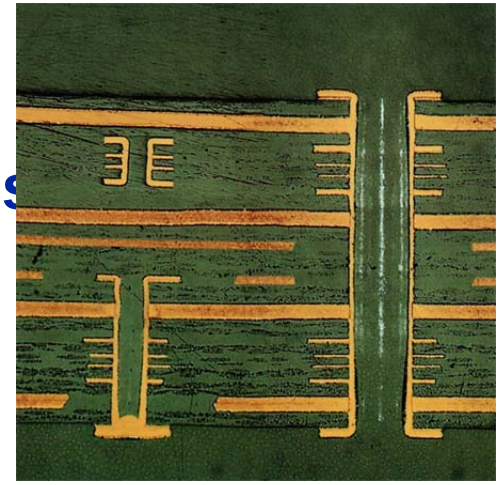

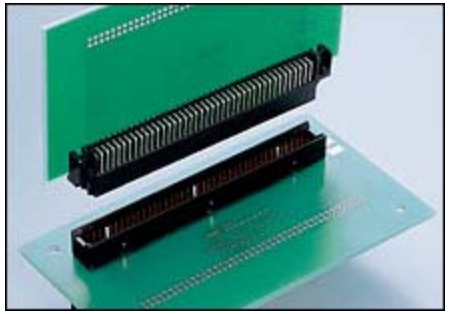

Line losses for 30” trace

Skin Effect
(3 dB/octave)

Dielectric Loss
(6 dB/octave)

Connectors, vias, etc.

S21=Pout/Pin for trace, in dB

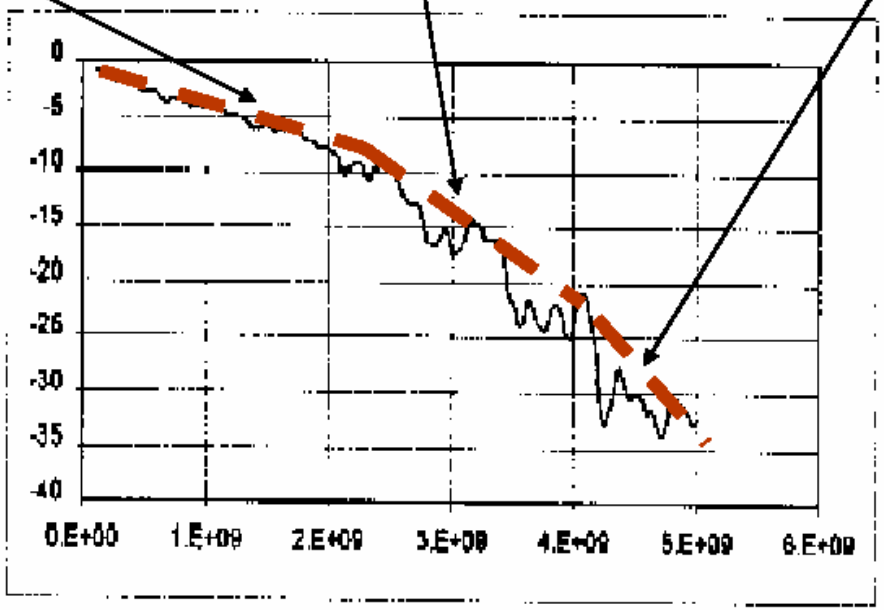

Figure 3 – Typical S21, transmission plot for serial link

d picture of what is happening in the frequency domain.

Crosstalk

- Worse than signal at above 10 - 18 GHz!

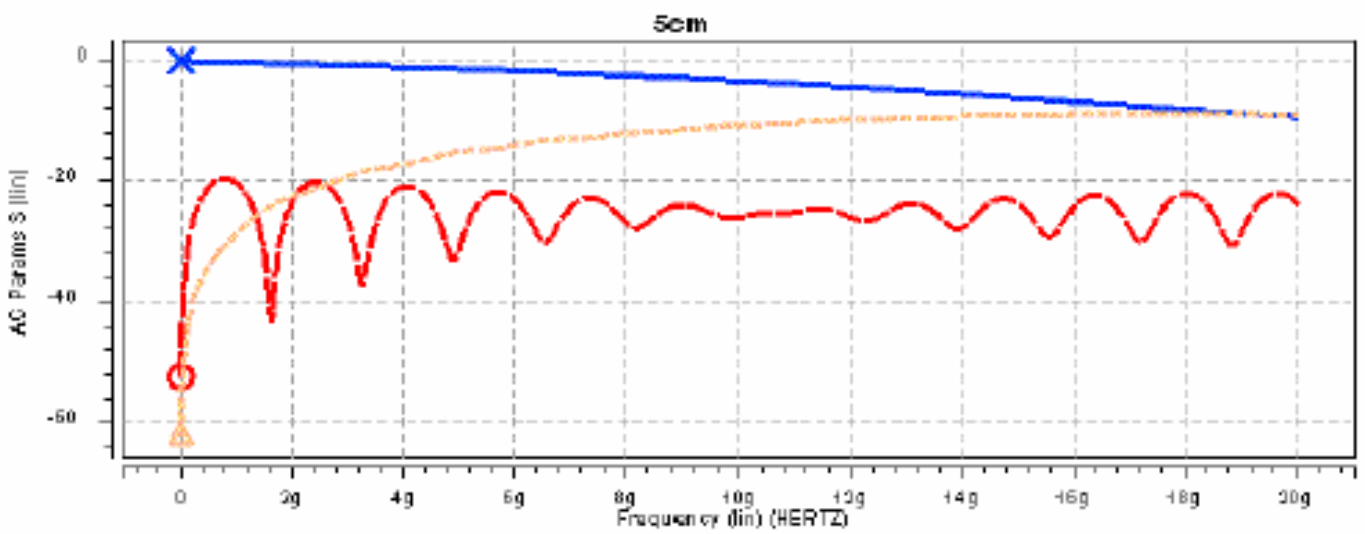

Signal
Far-end Xtalk

Near-end Xtalk

Solution Paths

- **Circuits and Signaling Schemes**
  - AC Coupled Interconnect
  - Technique to increase wire density
  - Other techniques being pursued industrially
- **Technology**
  - 3D Technologies
  - SiP and Silicon Circuit Board
- **Design**
  - Reducing need for memory Bandwidth

AC Coupled Interconnect

- **Circuit Approach only**
  - Using on-chip MIM capacitors
  - Or package buried caps
  - Benefits:
    - Low power; Low circuit area; ESD Protection; Straightforward Design For Test

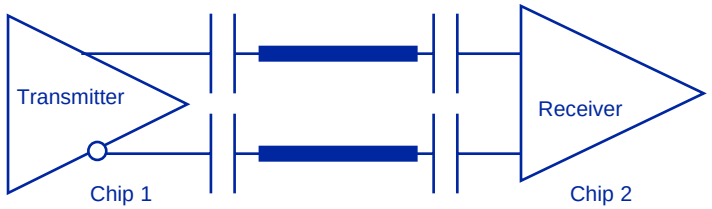

- **Package and Circuit Approach**
  - Build capacitors using chip-package interface
  - Additional Benefit:
    - High Density: 65-70 μm AC pad pitch

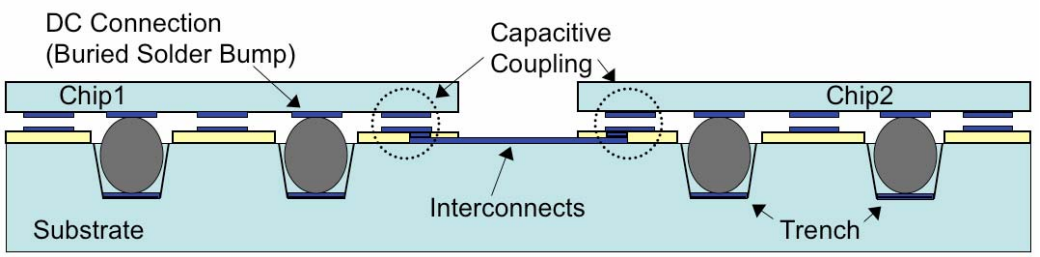

Benefits

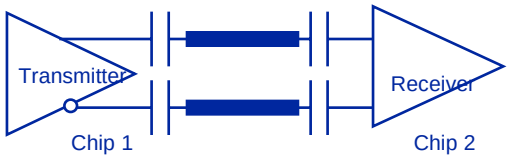

- **Circuit Benefits**
  - Capacitors on–chip, on-package or between
  - Power = 12 mW per channel @ 6 Gbps
    - About 3x less than conventional signalling
  - Circuit area ~ 5x less than conventional
    - Permits easier floorplanning
  - No ESD protection needed (unverified)
- **Package Structure Benefits**
  - Capacitors formed between chip and package
  - E.g. 4,800 power/ground, 4,200 signal on an 18x18 mm chip
    - High signal I/O and improved power/ground delivery

Pulse Signaling - Overview

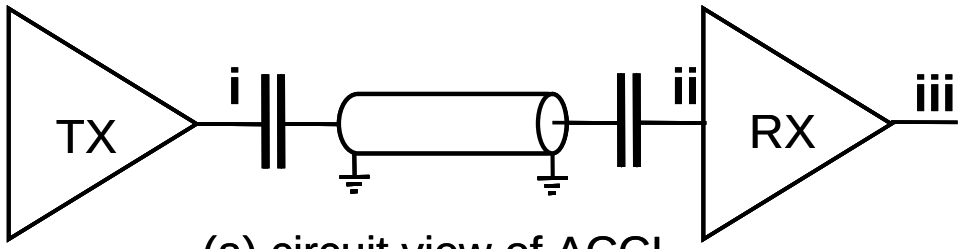

(a) circuit view of ACCI

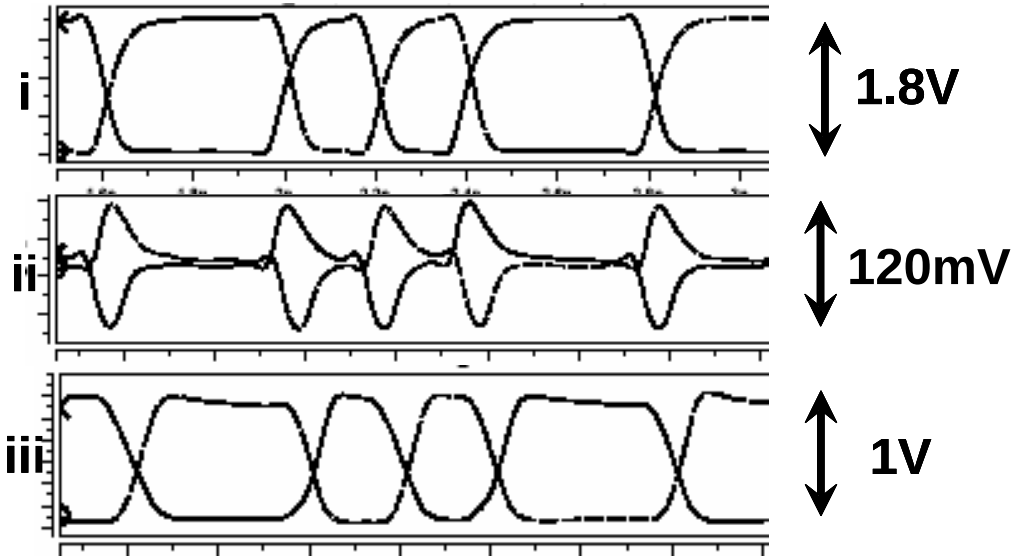

(b) pulse signaling waveforms

- 1st CC acts like differentiator
- 2nd CC acts like voltage divider
- NRZ data at TX output changes to pulses at RX input
- A simple RX perform equal and recovers NRZ data

Example of a traditional equalization

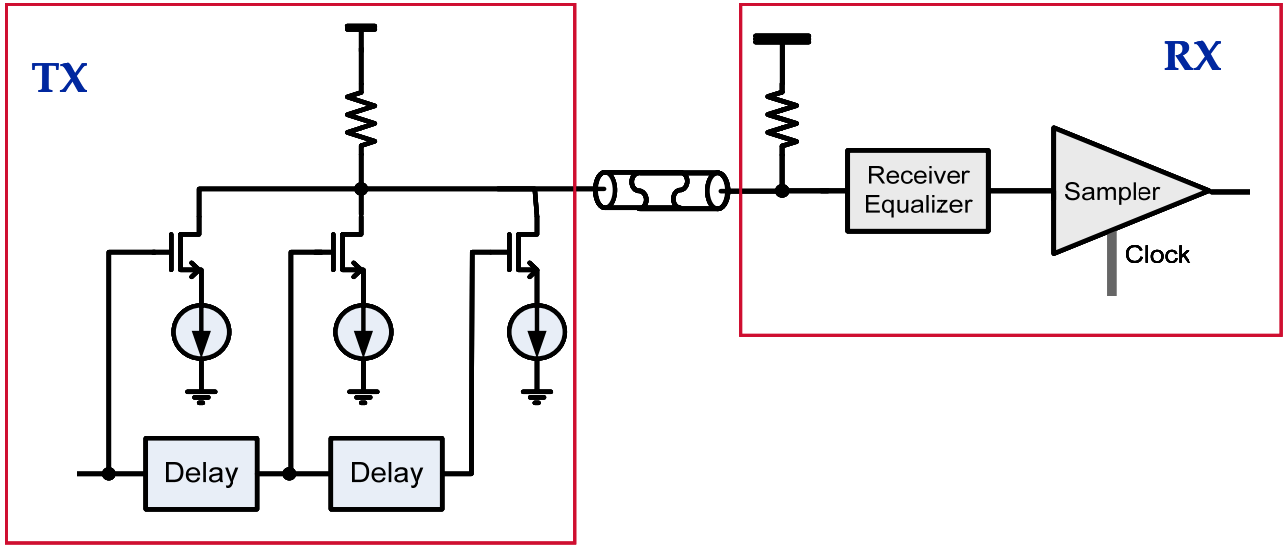

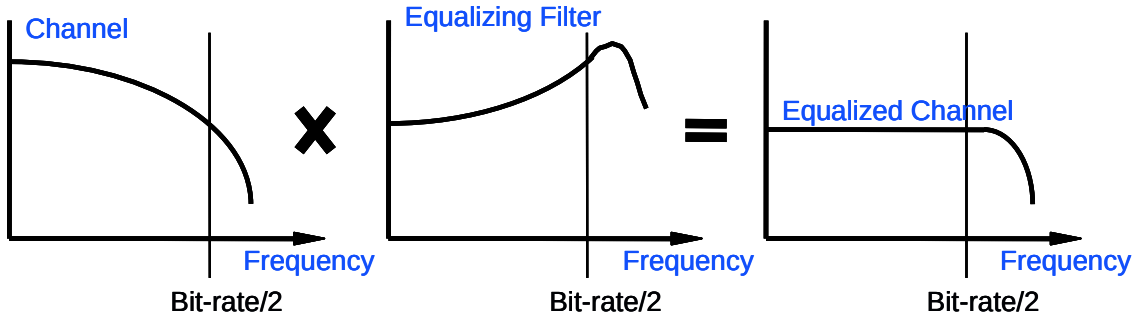

- Compensate high frequency loss on T-line: **complex**

Pulse signaling Equal – Freq Domain

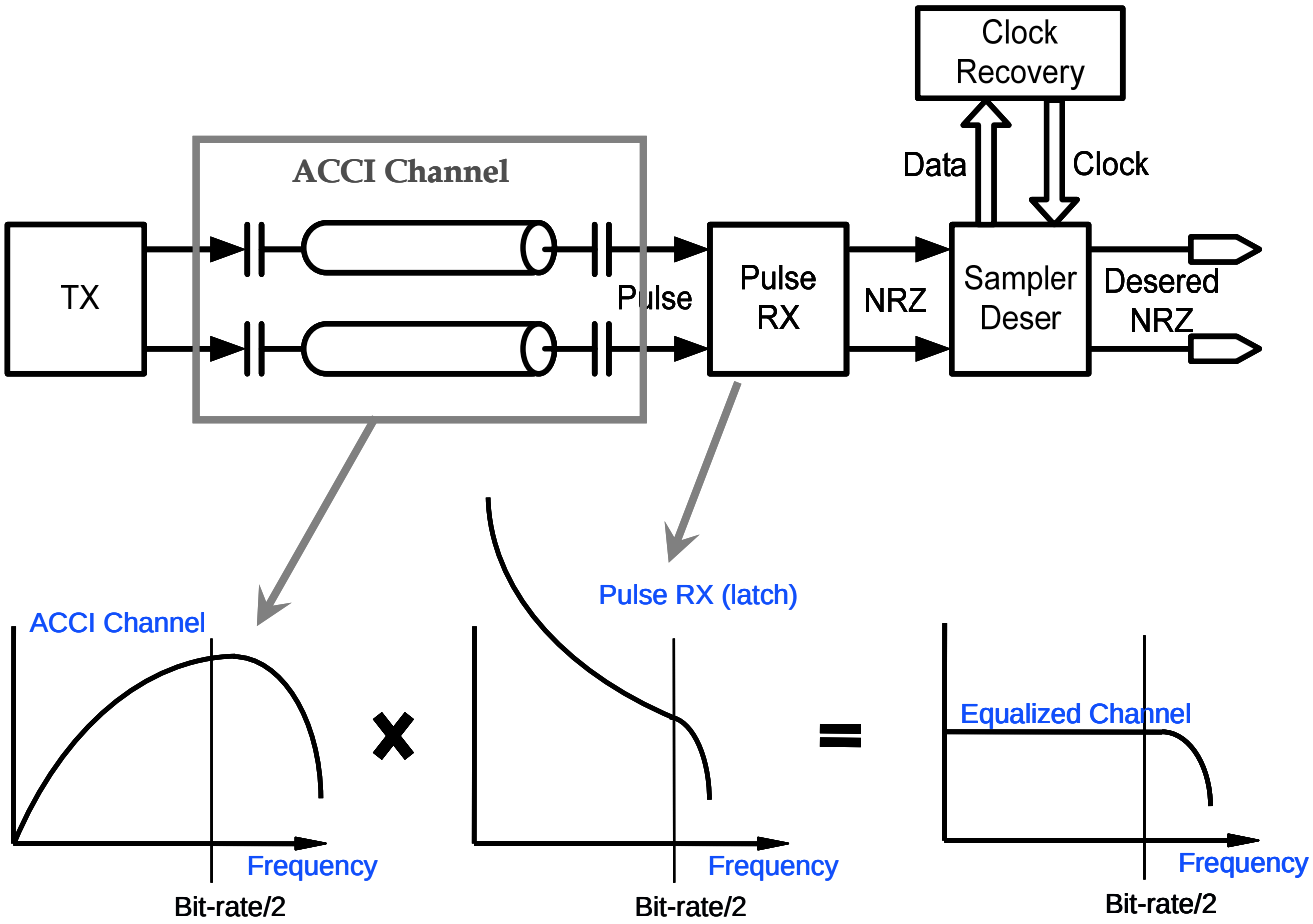

- Compensate low frequency loss on CC: **simple pulse RX**

Pulse Receiver Overview

- Blocks
- Voltage and timing margin budget

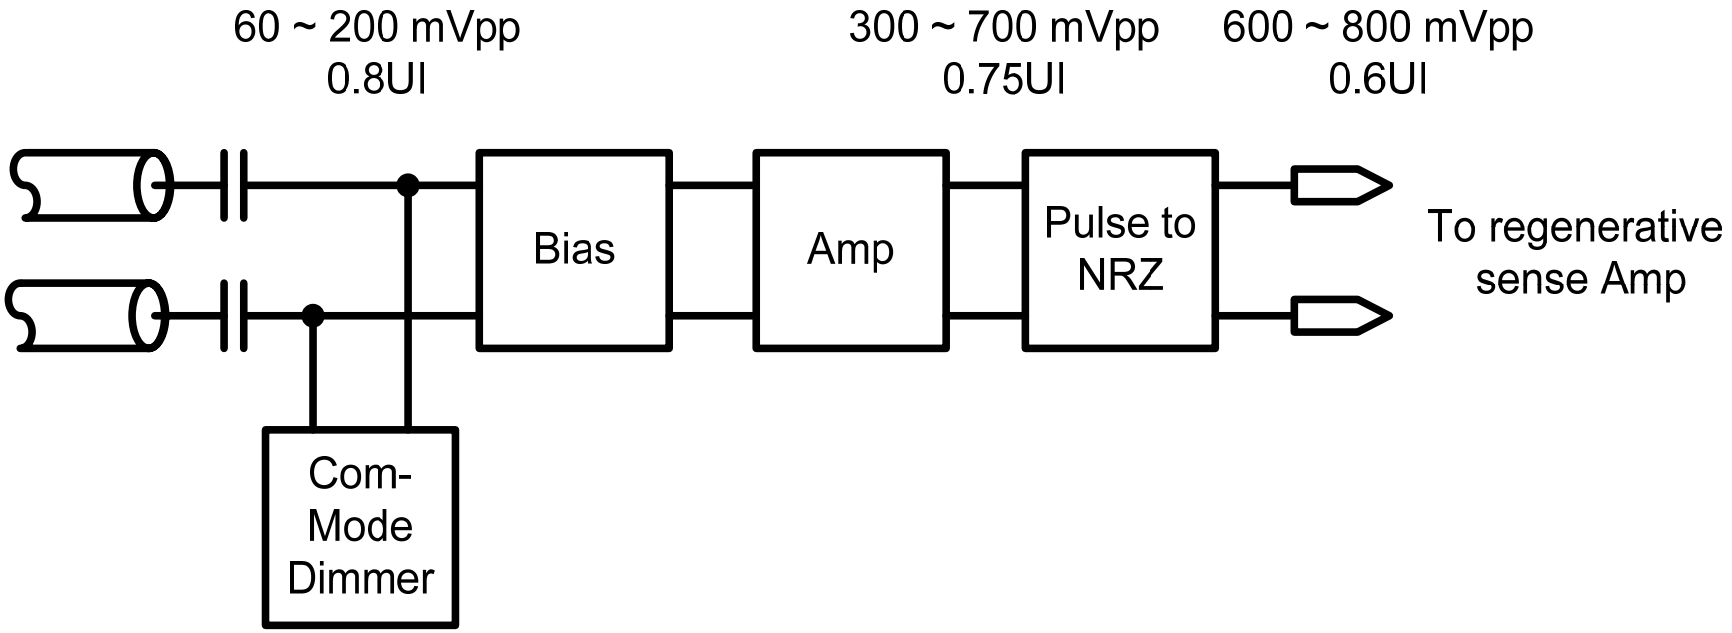

36 Gbps Circuit Demonstration

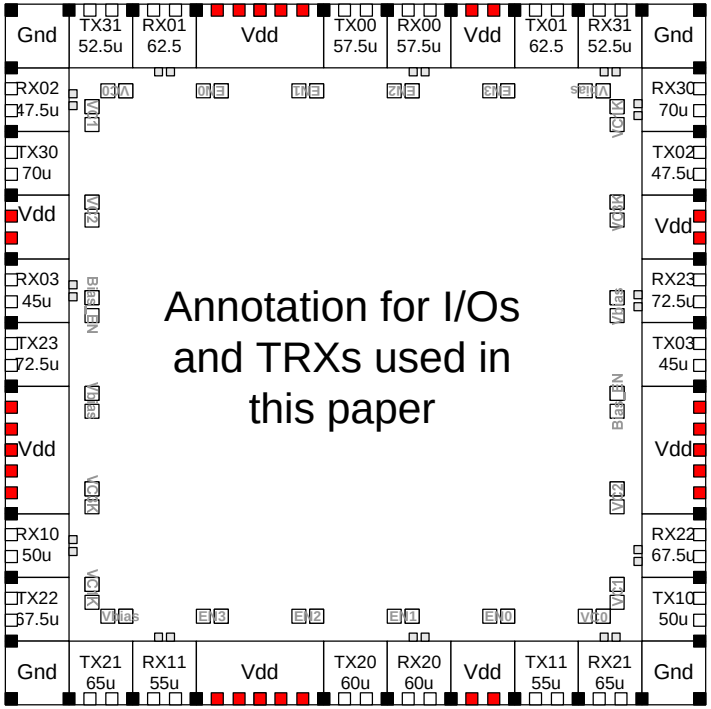

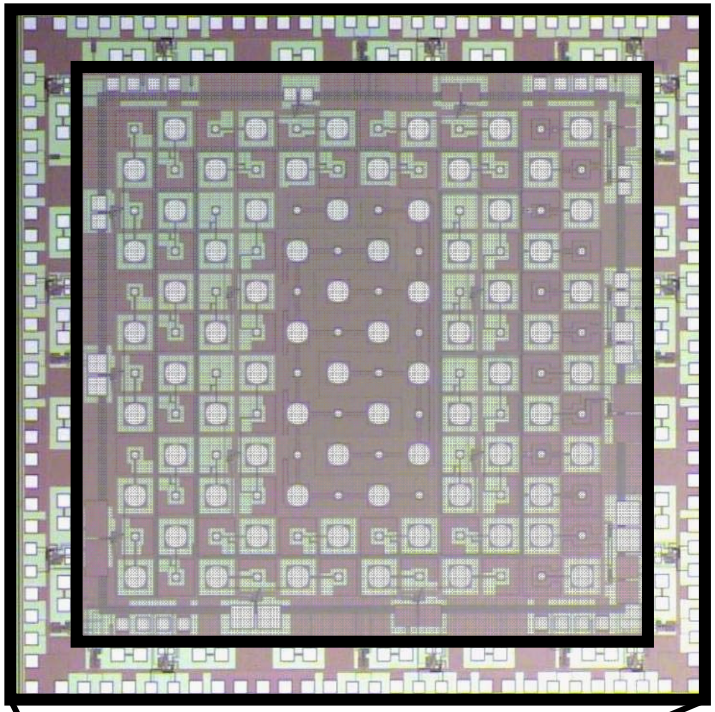

Die photo

Measured RX output at 6Gb/s operation

Differential signal

Single ended signals +

Single ended signals -

Differential Eye

Single ended Eyes +

Single ended Eyes -

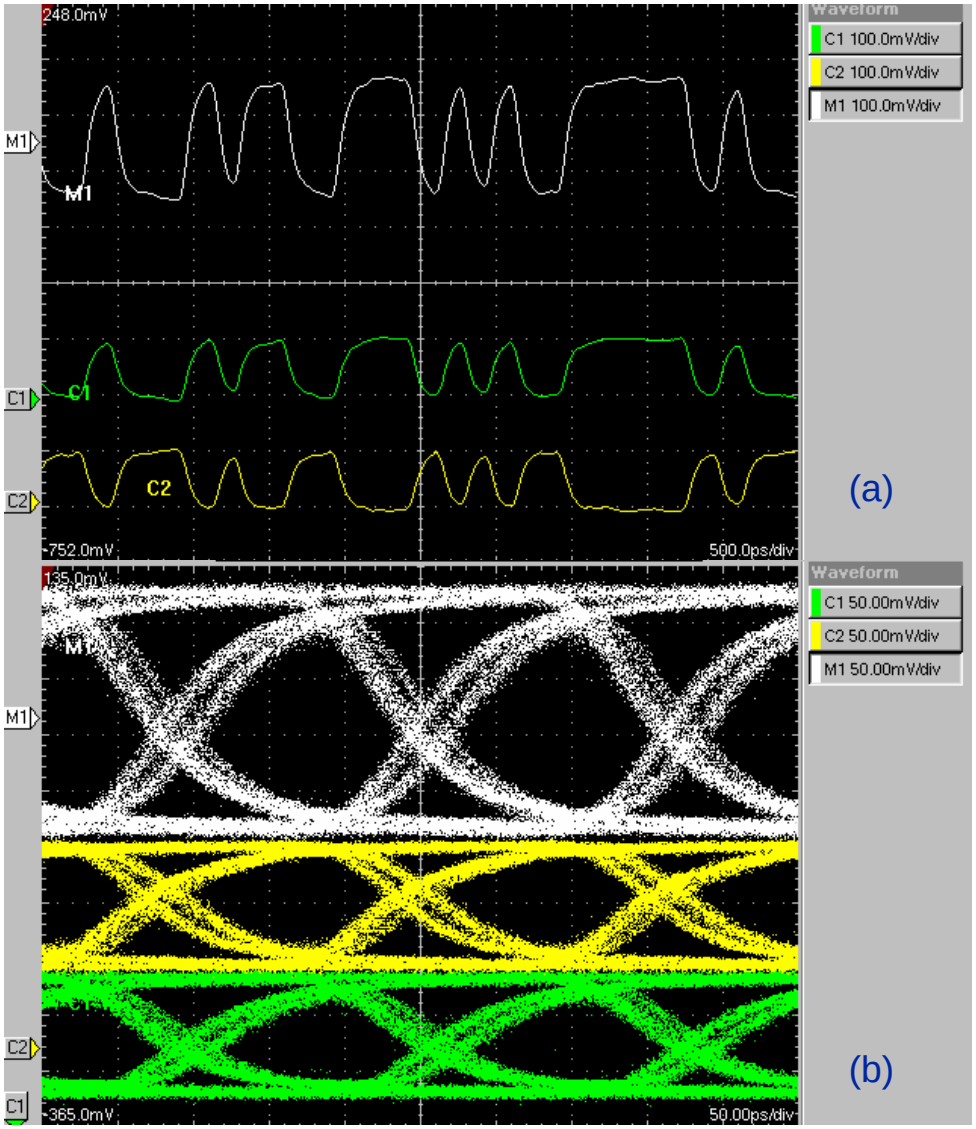

50% Power Savings

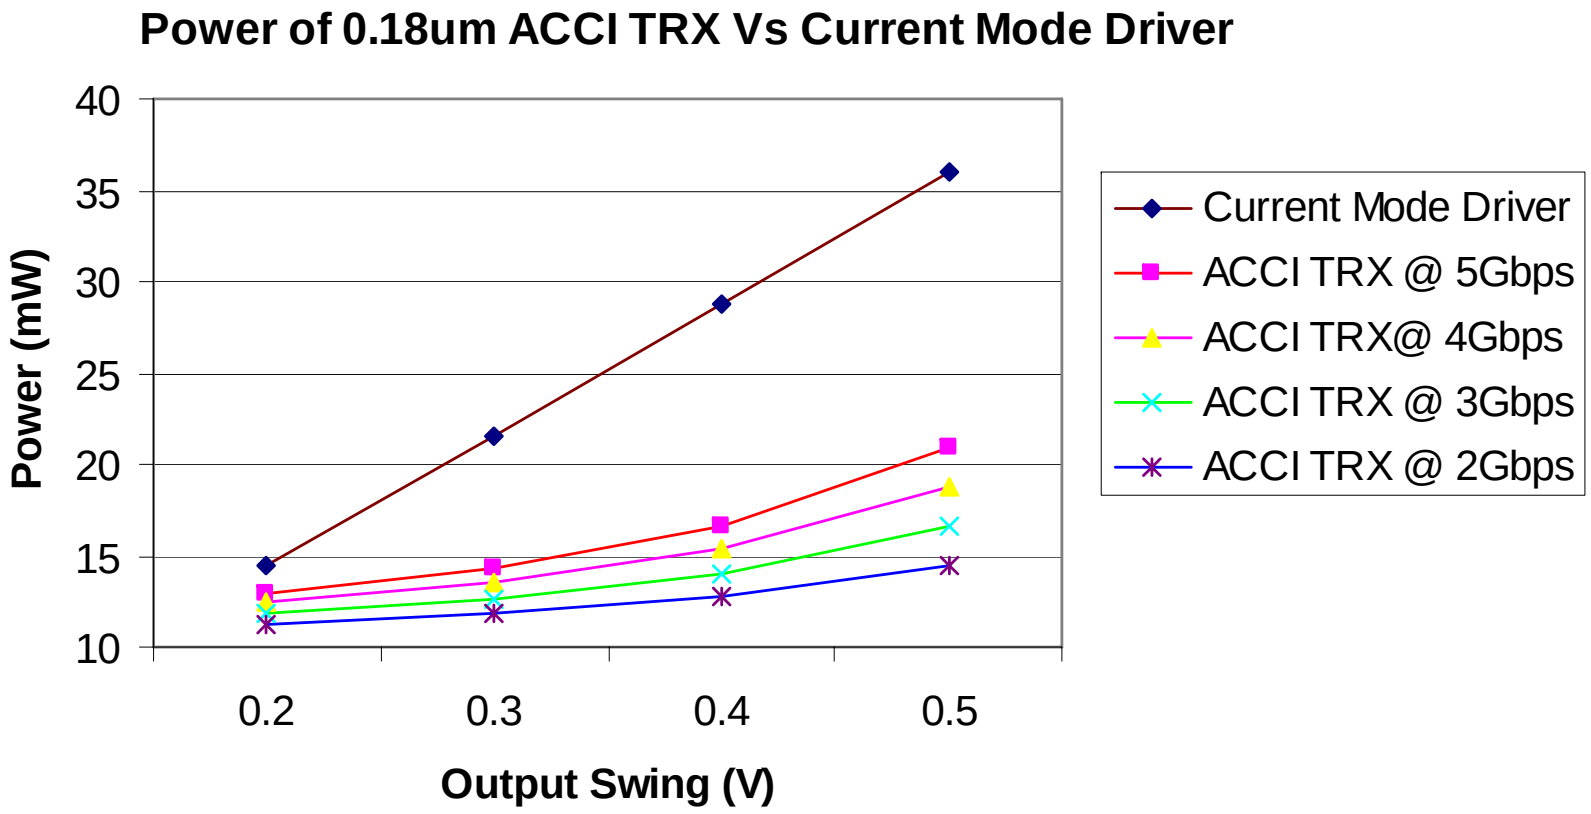

RX Input Pulses with 50ohm Stubs

4mil Stub

20mil Stub

40mil Stub

60mil Stub

80mil Stub

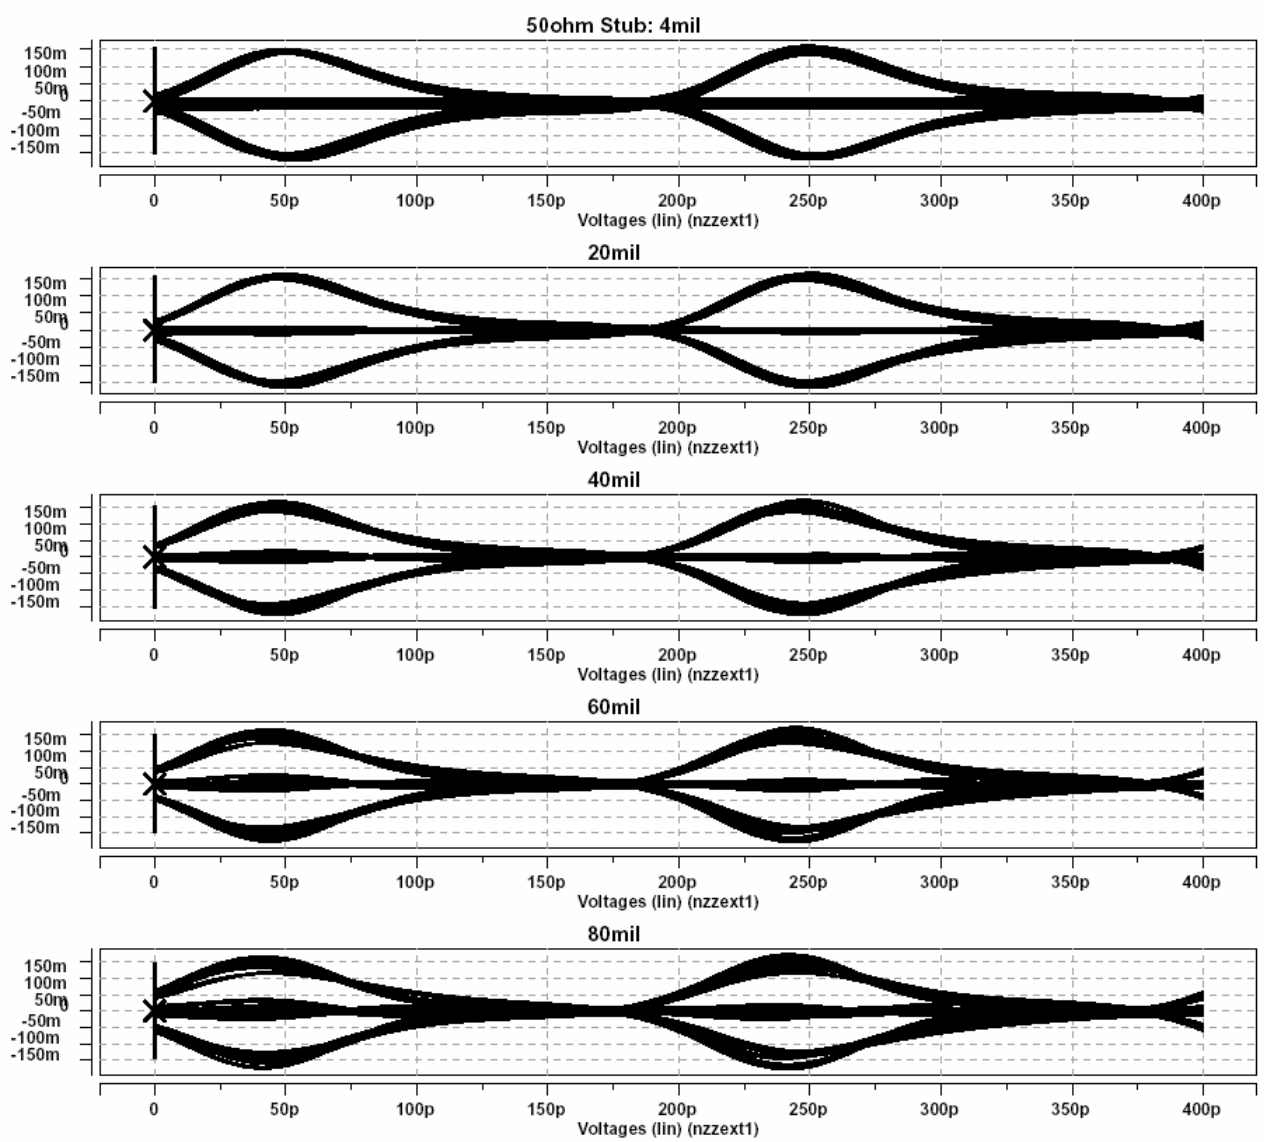

20 Gbps ACCI Channel

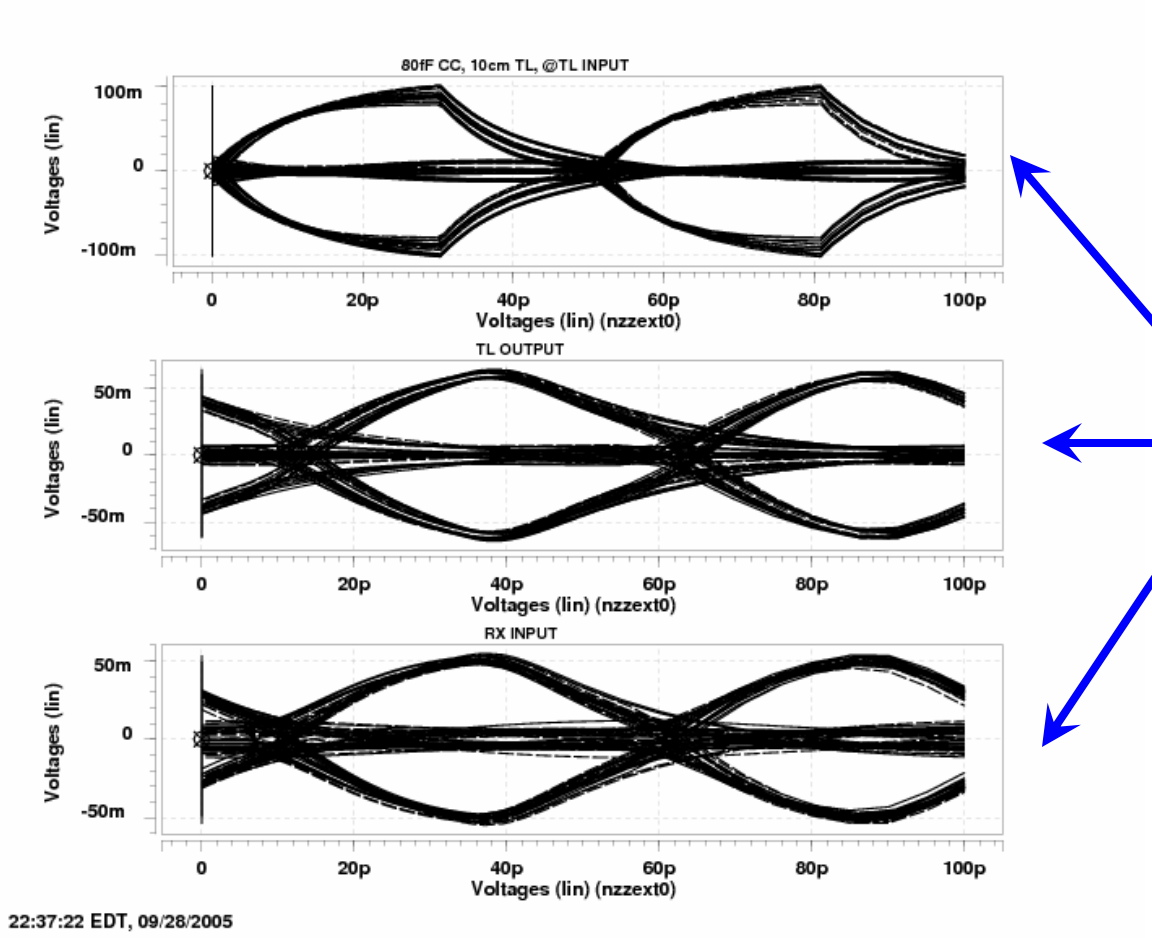

- 0.7 V Vdd
- 20Gb/s waves
- At T-line input
- At T-line output
- At RX input
- 80fF CC size
- 10cm T-line
- Swing limited at RX input

Inductive Coupling

- **Beyond scope of immediate program**
- **Main applications:**
  - High density, low-cost Connectors and Sockets
  - 3D ICs
- **Circuits:**

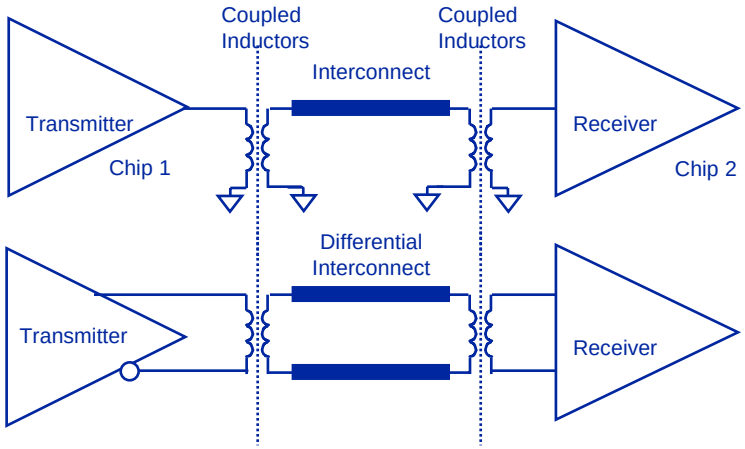

Inductive Coupling in 3D ICs

- Experiment performed:

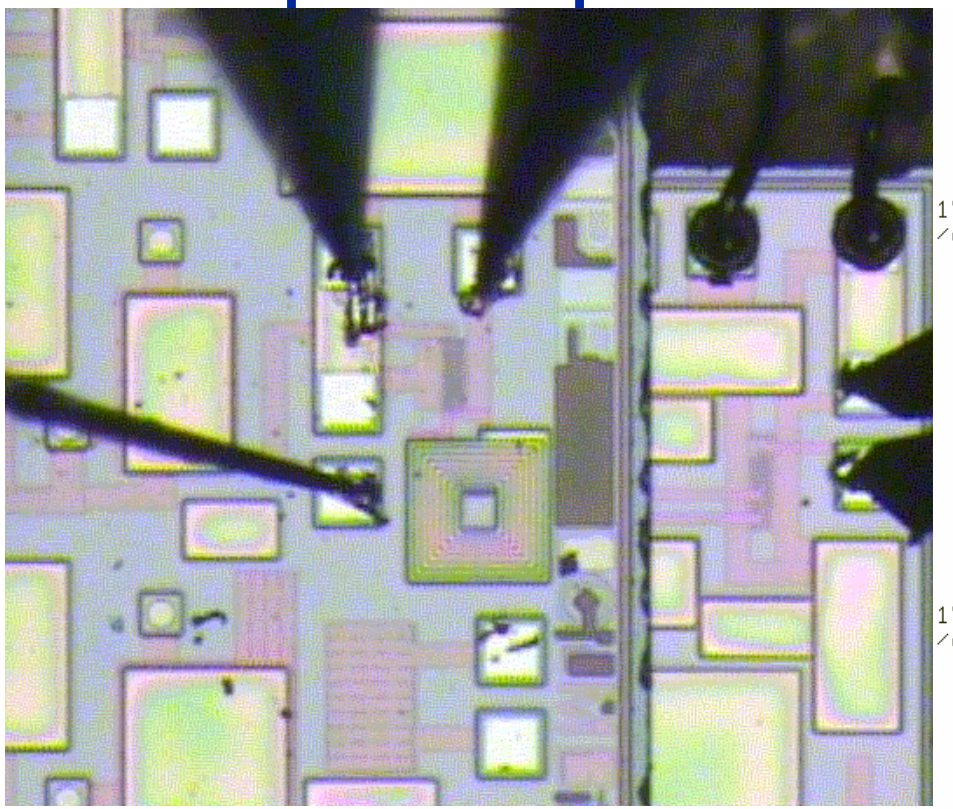

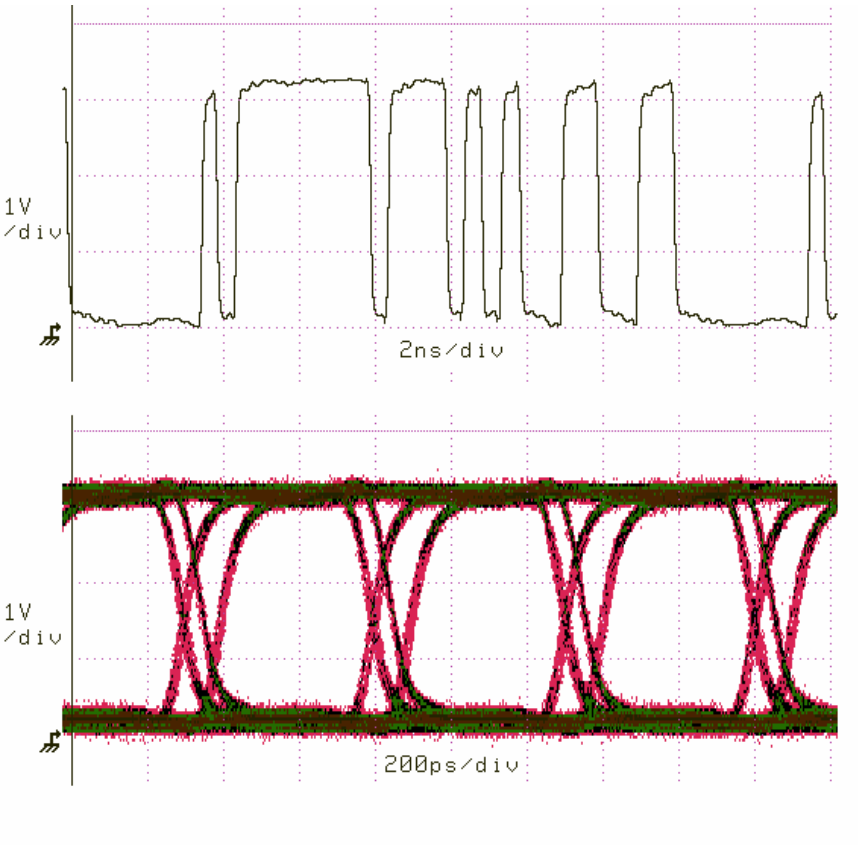

Connector Proof of Principle

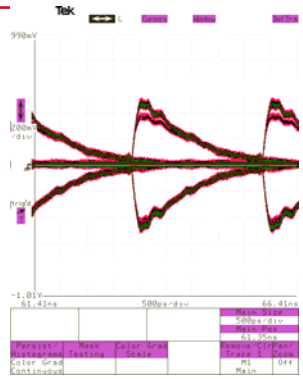

- PCB mockup

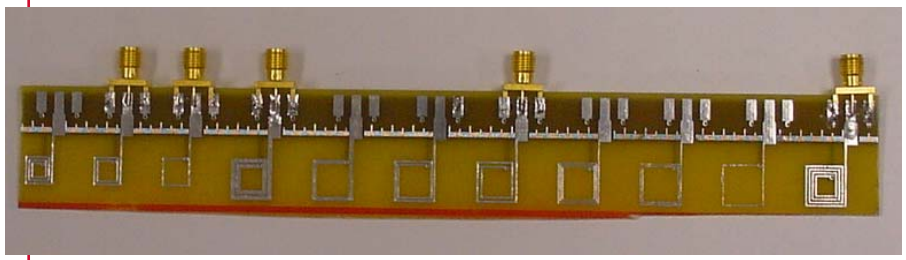

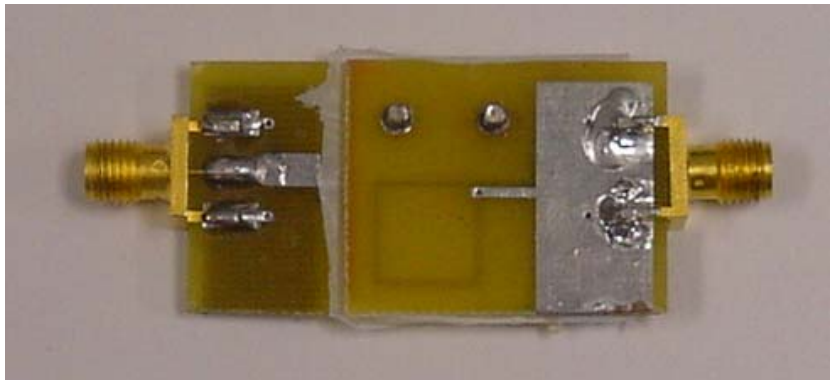

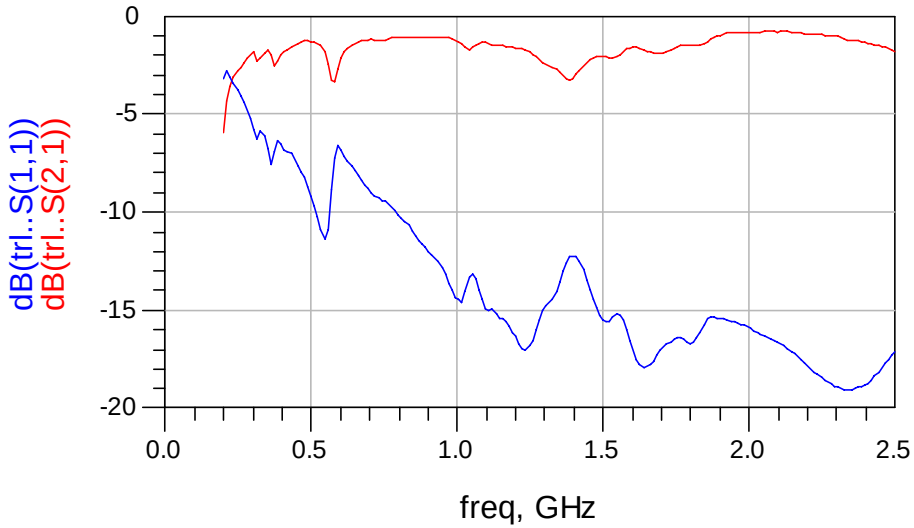

Increasing Interconnect Density

- **Further Out: Problem Statement:**
  - High density chip I/O mainly benefits power/ground system
  - To impact system bandwidth, need to increase wire density and Gbps/cm-cross-section in PCB
- **New Idea**
  - Eliminates crosstalk
  - Increases symbol rate without an increase in GHz
  - Increases BW/X-section 3x or more

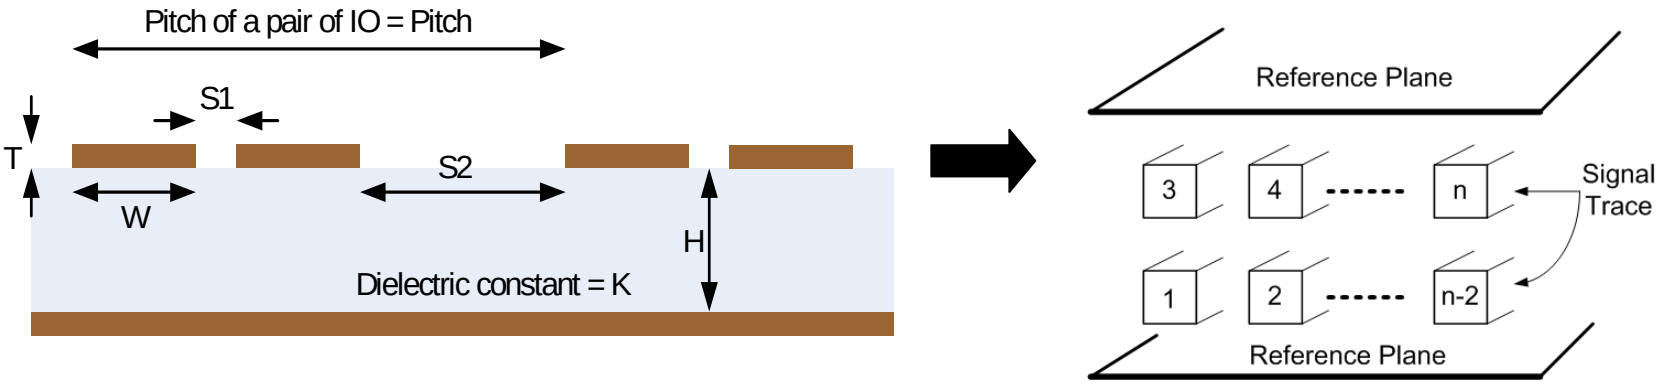

3D Technologies

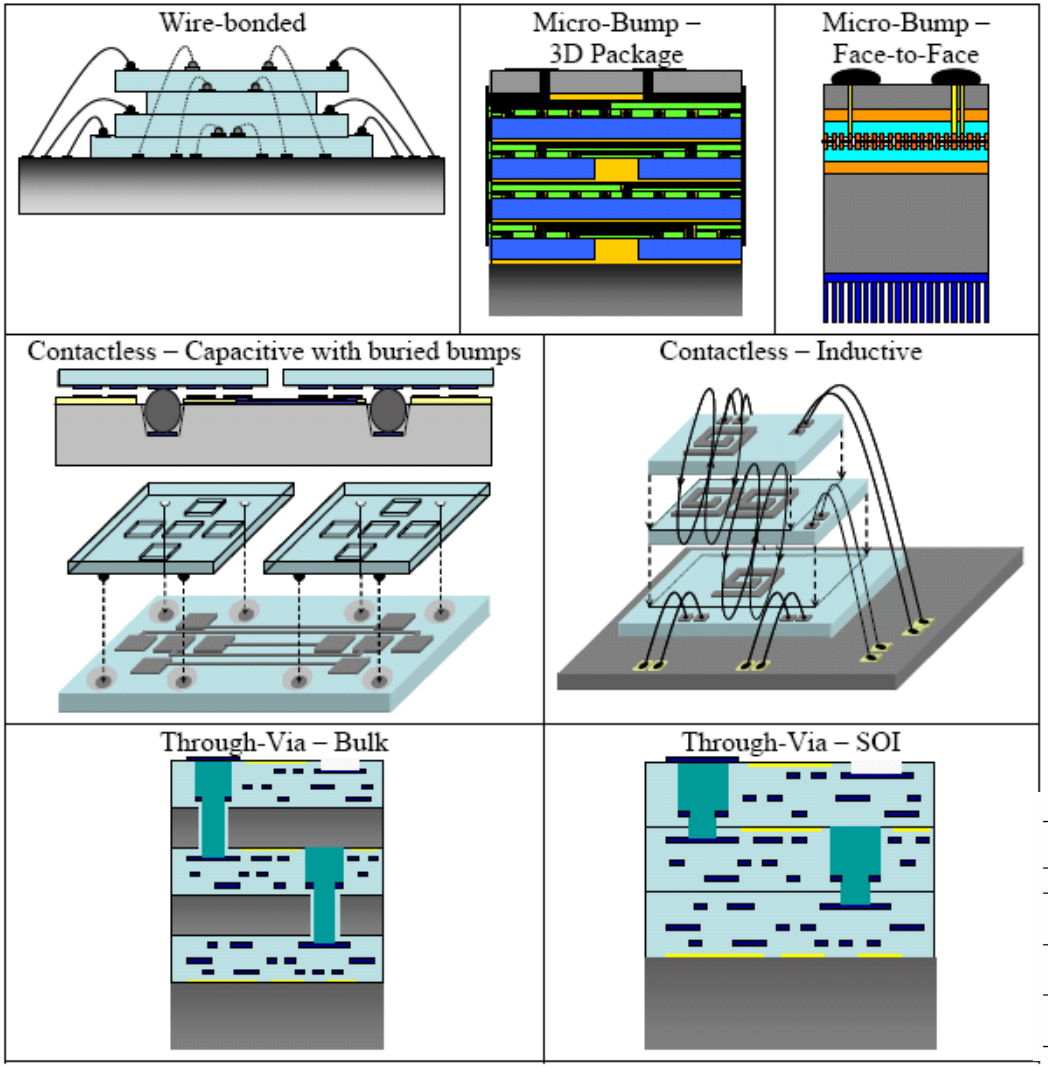

Table 1. Comparison of Vertical Interconnect Technologies

		Assembly	Tier limit	Vertical Pitch	Chip Layer Resources
Wire-bonded		Die	~5	35-100µm	All
Micro-bump	3D Package	Die	heat	25-50µm	Top 1-2
	Face-to-face	Die	2	10-100µm	Top 1-2
Contactless	Capacitive	Die	2	50-200µm	Top
	Inductive	Die	heat	50-150µm	Top 1-2
Through-Via	Bulk	Wafer	heat, yield	50µm	All + Top
	SOI	Wafer	heat, yield	5µm	All + Top

3D IC and 3D Packaging Technology

- **3D technologies we are working with**

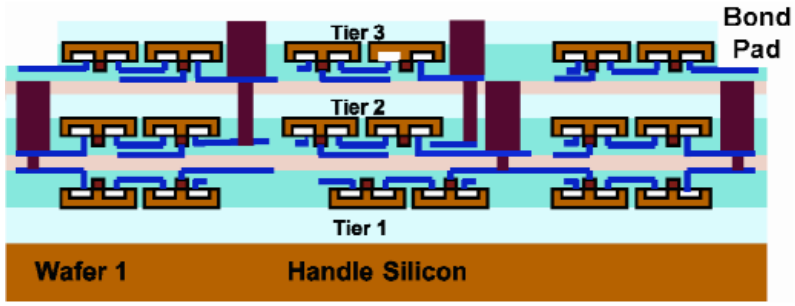

c/- MIT LL

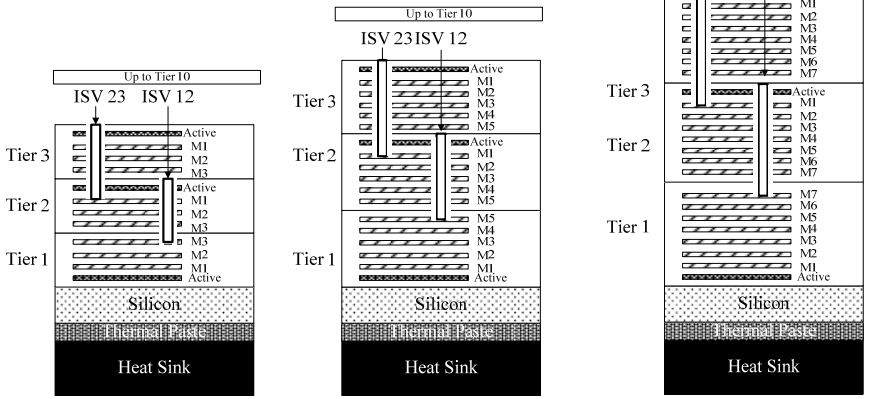

c/- Hao, Davis, NCSU

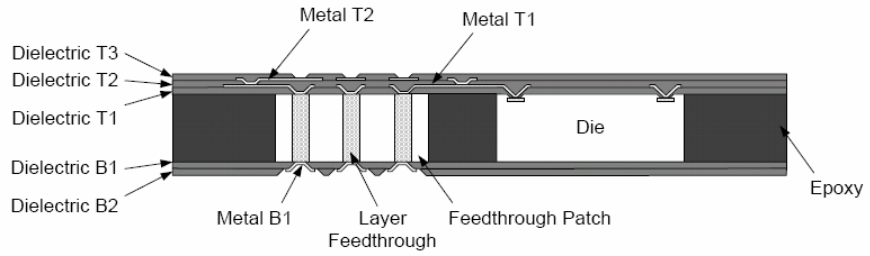

Silicon "PCB"

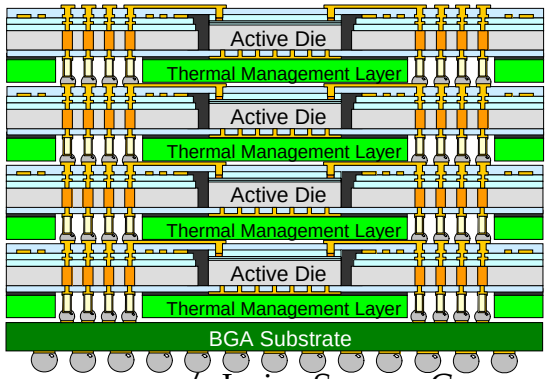

c/- Irvine Sensors Corp

2-core Processor Case Study

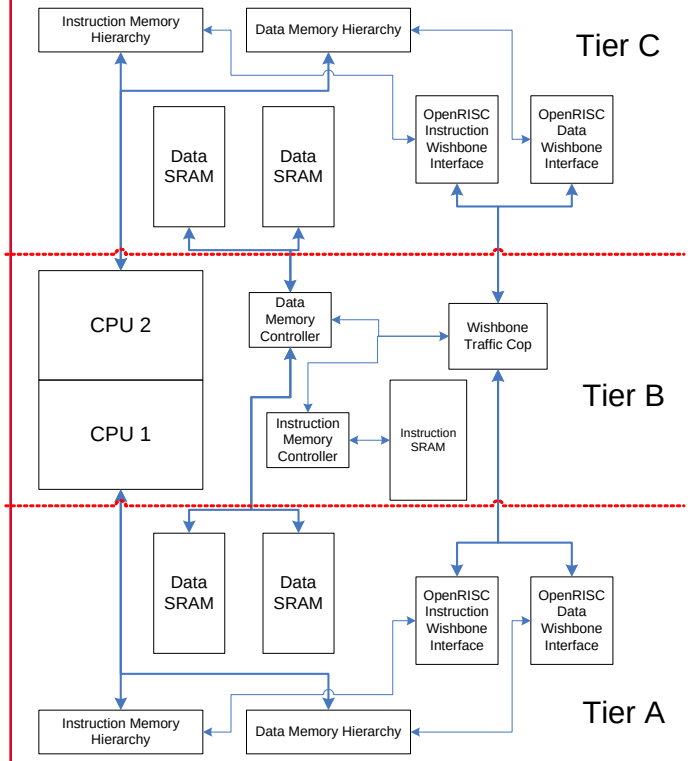

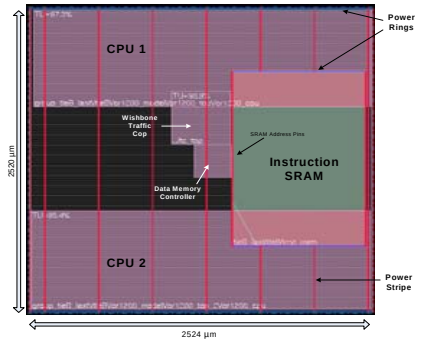

Tier B

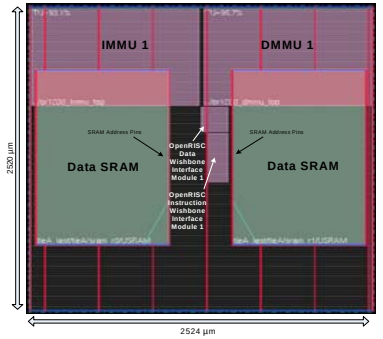

Tier A

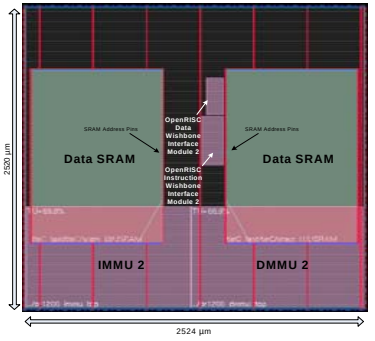

Tier C

		2D	3D	% Reduction
Critical Path		**25.13 ns**	**17.87 ns**	**28.89**
Data Memory	**To SRAM**	**33.07 ns**	**23.0 ns**	**30.45**
	From SRAM	**36.6 ns**	**24.86 ns**	**32.07**
Instruction Memory	**To SRAM**	**34.01 ns**	**23.27 ns**	**31.58**
	From SRAM	**37.82 ns**	**24.01 ns**	**36.51**

c/- Schoenfloss, Davis, NCSU

Hardware Algorithmic Approaches

- **Key Enablers:**
  - Best exploiting DRAM architecture
    - E.g. DDR2-400
      - Burst mode bandwidth : 400 Mbps
      - Random mode access : 16 Mbps

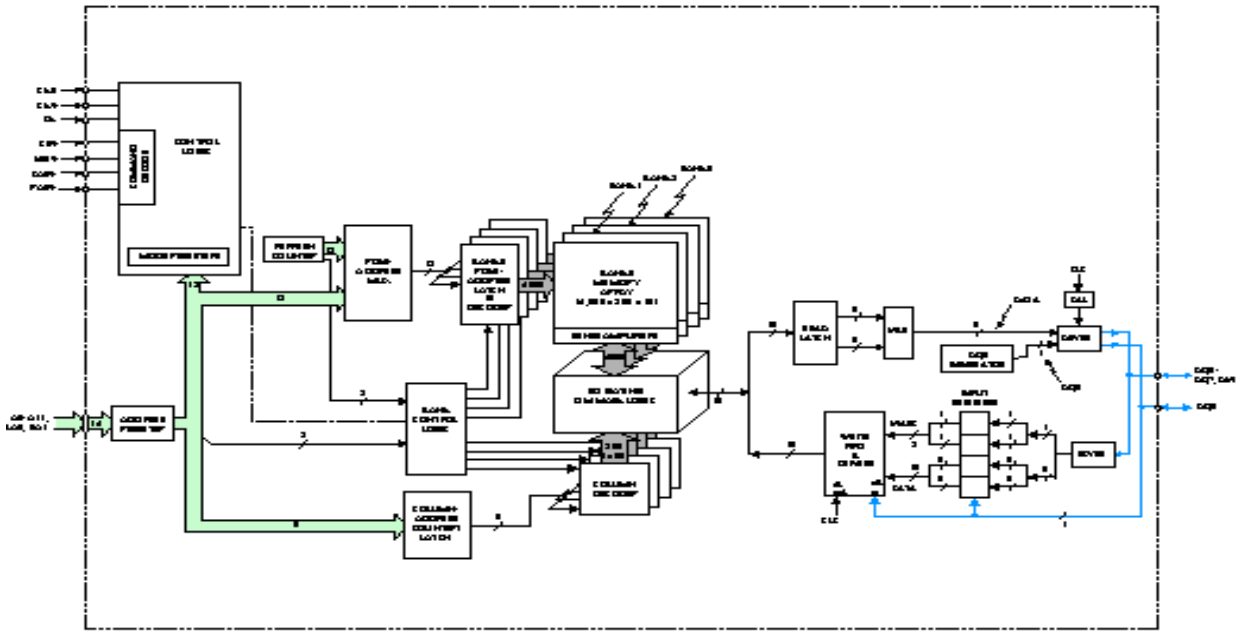

SAR DSP Case Study

Performance: 1 Million point, 1 ms FFT for future radar systems

Solution	Size	Power
8-tier 3D MINTs with COTS components (FPGA/memory) 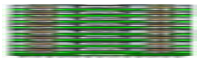 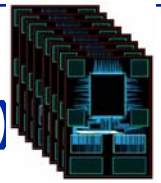	**35x25 x7mm**	**70 W**
4-tier 3D MINTS with ASICs + COTS memories 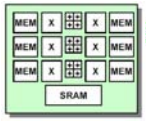 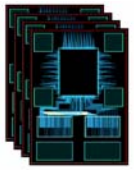	**35x25 x3 mm**	**50 W**
12-tier 3D IC with ASICs + custom memory 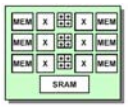 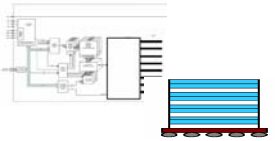	**16x12 x3 mm**	**6 W**
Off-the-shelf solution using commercial DSPs, memories and boards (all COTS)	**540x540 X10 mm**	**80 W**

Hardware Algorithms

- Hardware algorithms can be used to reduce required bandwidth to DRAM
- E.g. FFT: Organized addresses to maximize use of burst mode ➔ Increased performance over 50x

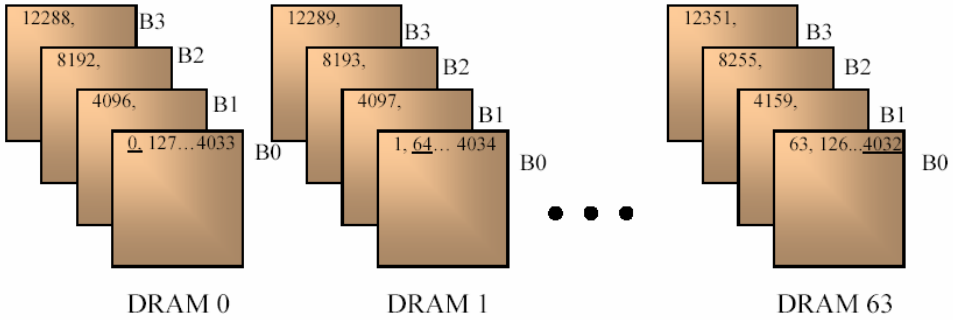

Reads:
Row # = ⌊FFT#/4⌋
Bank # = FFT# % 4

Writes:
Row # = ⌊FFT#/256⌋
Bank # =(⌊FFT#/64⌋) % 4

Where FFT# = ⌊index/64⌋

Optimal chip-package codesign for high-performance DSP
Mehrotra, P.; Rao, V.; Conte, T.M.; Franzon, P.D.;
Advanced Packaging, IEEE Transactions on [see also Components, Packaging and Manufacturing Technology, Part B: Advanced Packaging, IEEE Transactions on]
Volume 28, Issue 2, May 2005 Page(s):288 - 297

Networking

- **Scheme to reduce DRAM requirements for Networking**
  - IP Forwarding, Firewalls, Diffserv
  - Key: Store compressed search data in Trie
  - 28M lookups/s in 1 sq.mm. of 0.13um Silicon
  - Trie configured by instruction extension

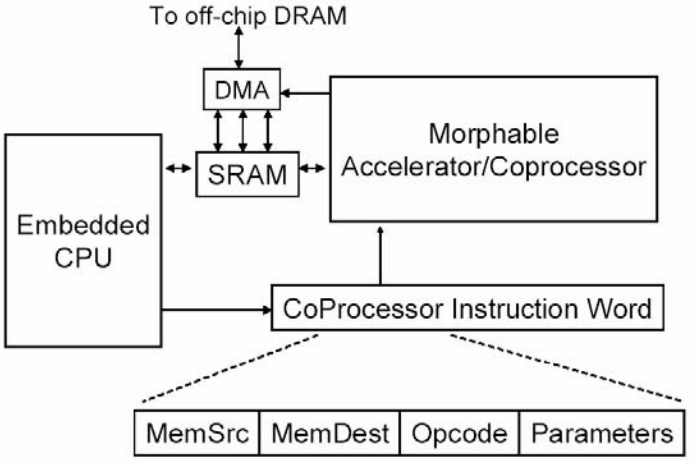

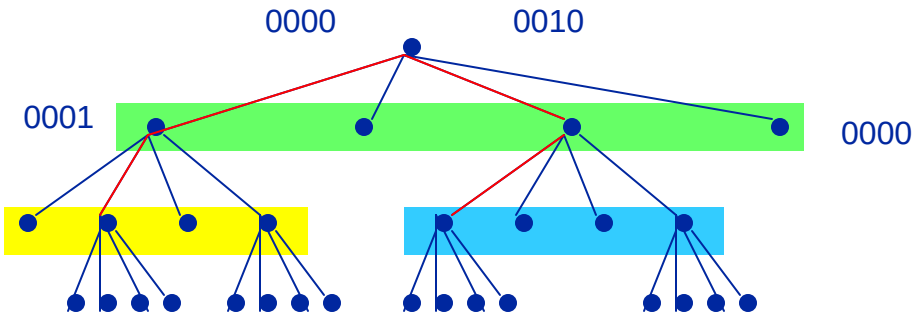

M. Yadav, P. Hamilton, R. Sears, Y. Viniotis, T. Conte, P.D. Franzon,
"A configurable classification engine for polymorphous chip architecture,"
ACM BEACON Workshop, Boston, OCT. 2004

Novel hardware architecture for fast address lookups
Mehrotra, P.; Franzon, P.D.;
Communications Magazine, IEEE
Volume 40, Issue 11, Nov. 2002 Page(s):66 - 71

Speech Recognition

- Data organization techniques to reduce DRAM requirements
- Leads to two-chip solution

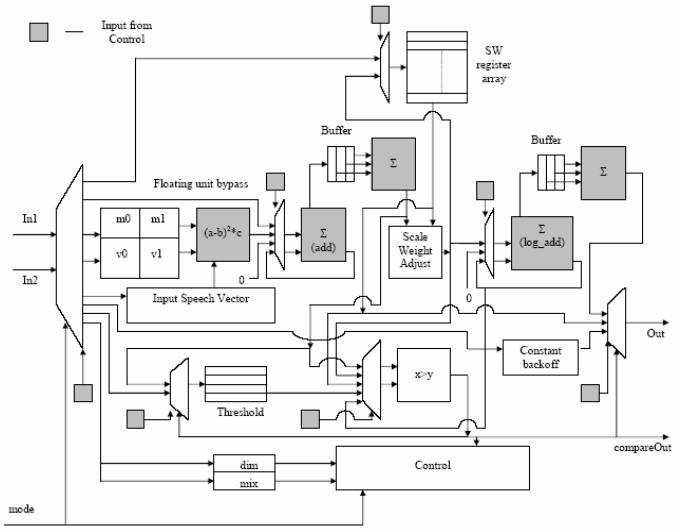

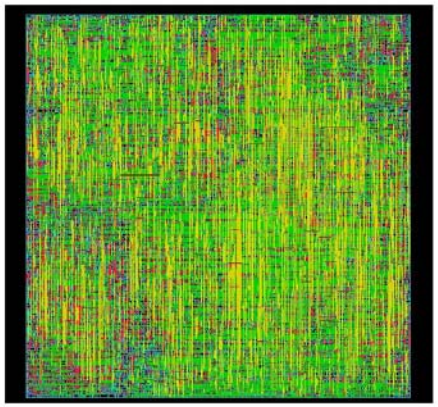

Conclusions

- **Memory bandwidth scales aggressively with Moores Law**
  - Leads to 1 TB/s by 2010
  - Drivers:
    - MultiCore; Graphics; Networking; Cognitive tasks
  - ➔ 12 – 40 W just for I/O!
- **Solution space**
  - Low-power, high-bandwidth, crosstalk-less I/O
  - Hardware algorithms that minimize DRAM usage

Acknowledgements

- **Funding Sources:**

- **Colleagues and students:**
  - Rhett Davis, John Wilson, Steve Lipa, Lei Luo, Monther Al Dwairi,