



# Multiplying Moore's Law with Proximity Communication

**Robert Drost, Ph.D.**  
**Director and Distinguished Engineer**  
Sun Microsystems Laboratories



**2007**  
**Sun Labs**  
**Open House**



# Outline

- The Bandwidth Motivation
- Proximity Communication Technology
- Multiplying Moore's Law

# The Team

## VLSI Research Group at Sun Labs


Igor Benko, Alex Chow, Wes Clark, Bill Coates, Robert Drost, Jo Ebergen, Scott Fairbanks, Jonathan Gainsley, Gilda Garreton, Yaeko Hirotsuka, Ron Ho, David Hopkins, Ian Jones, Russell Kao, Jon Lexau, Dimitri Nadezhin, Tarik Ono, Steve Rubin, Jeff Rulifson, Justin Schauer, Ivan Sutherland, and friends: David Harris, Mark Greenstreet, Ken Yang

## And many others at Sun



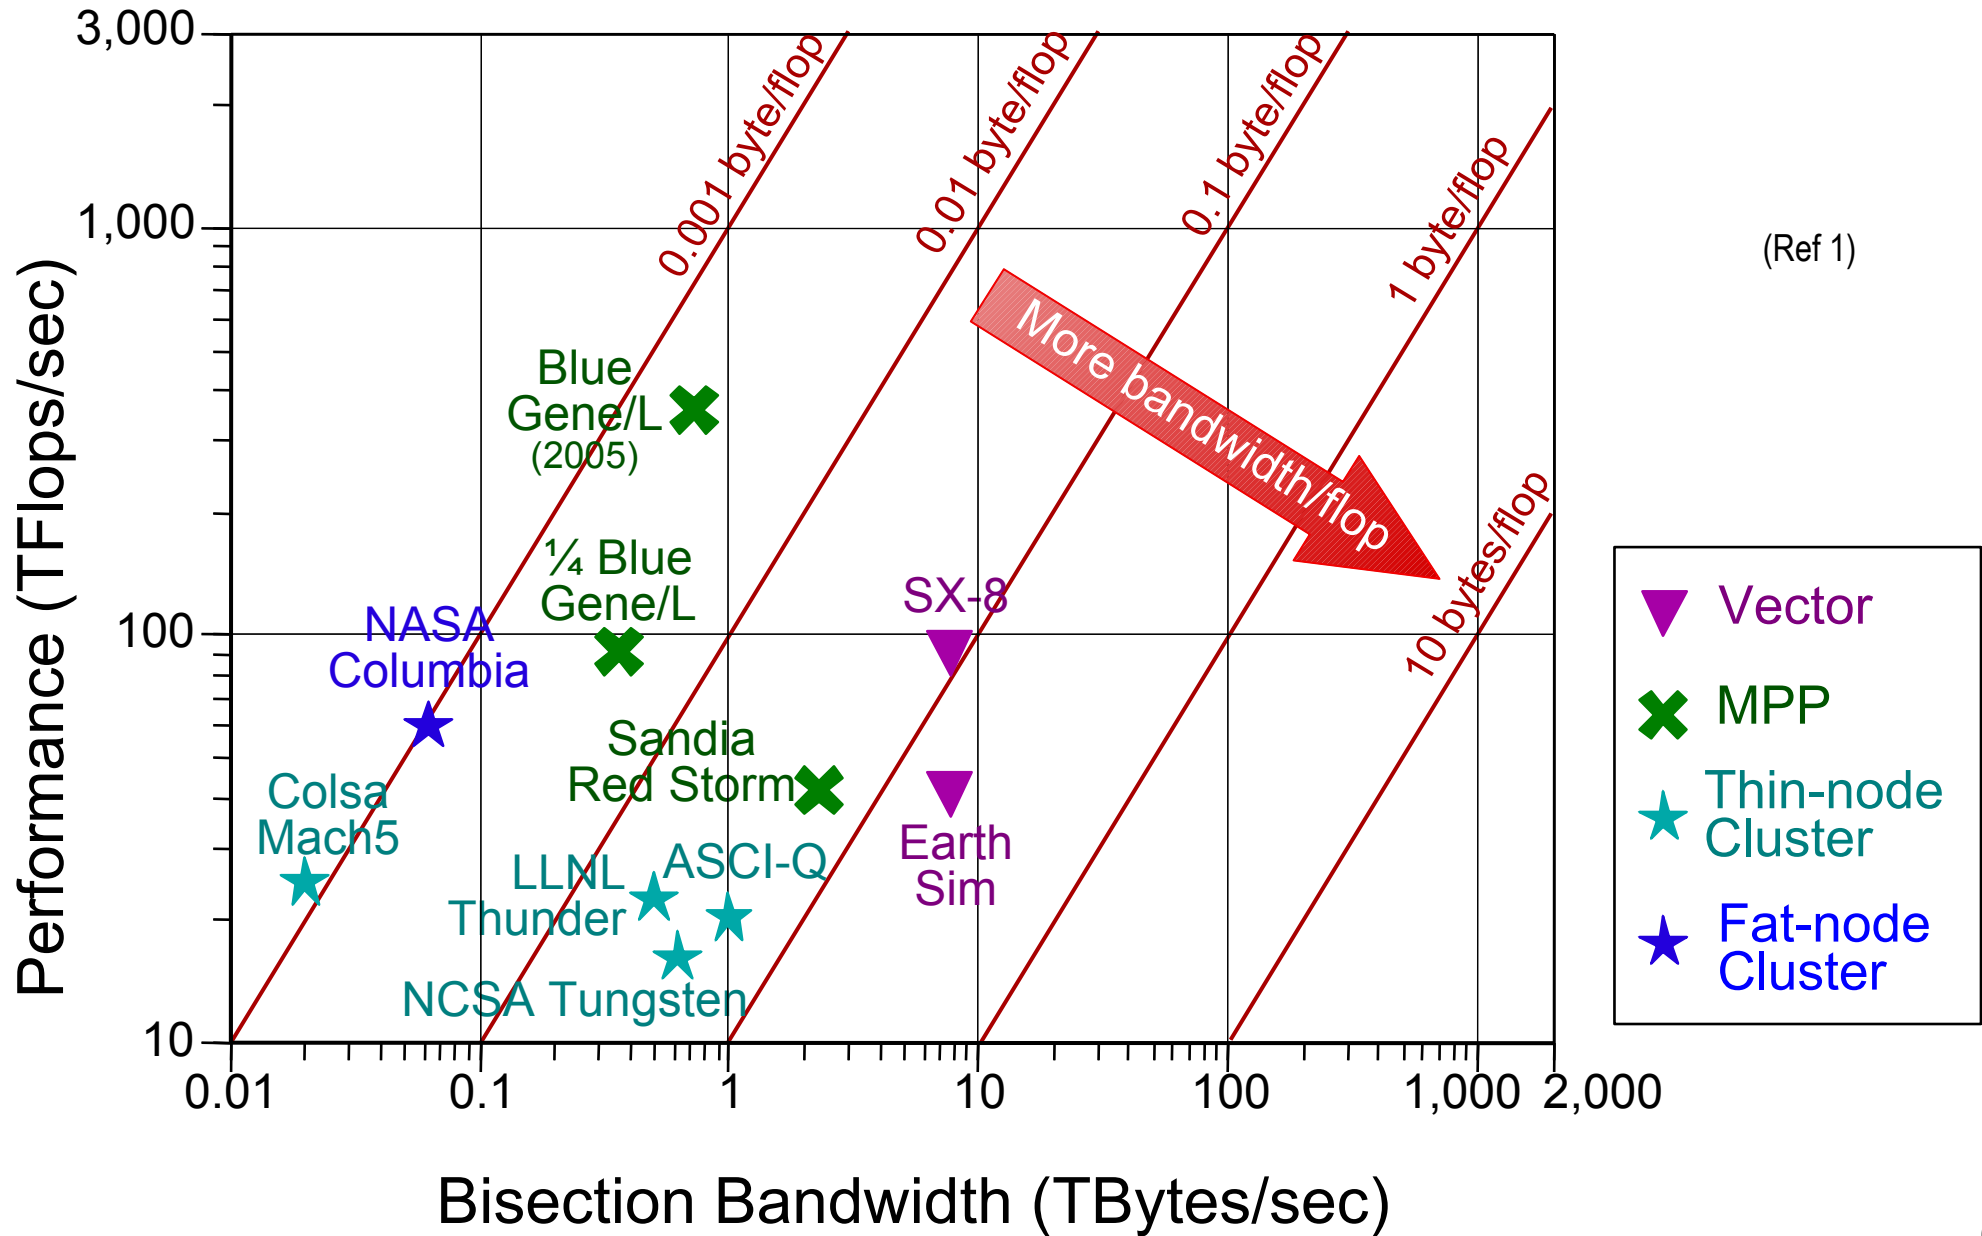
**Why do we want more off-chip bandwidth anyway?**

# Motivation: CPU vs. DRAM




J.L. Hennessy and D.A. Patterson, Computer Organization and Design, 2nd ed.

# Motivation: BBW vs. Flops



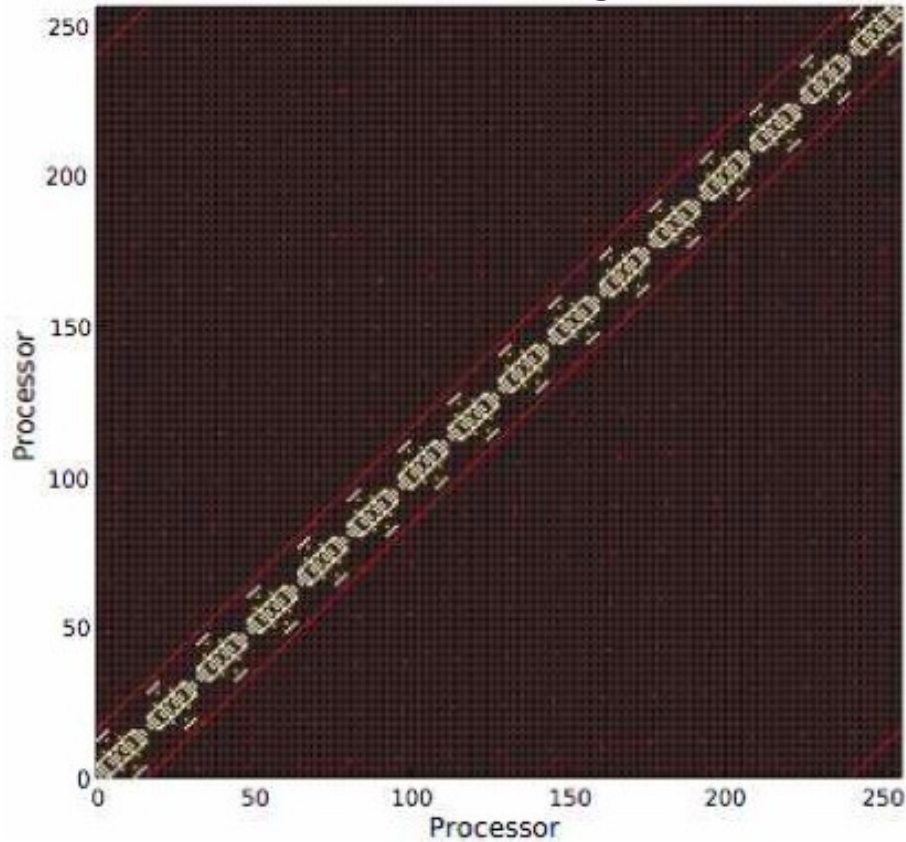
# Bandwidth versus Memory Capacity




(Ref 1)

# Motivation: Lack of Data Locality

## Dense Linear Algebra



## 3D FFT




(Ref 2)



Black=no processor pair communication  
 White=Heavy processor pair communication

# Proximity Communication Technology

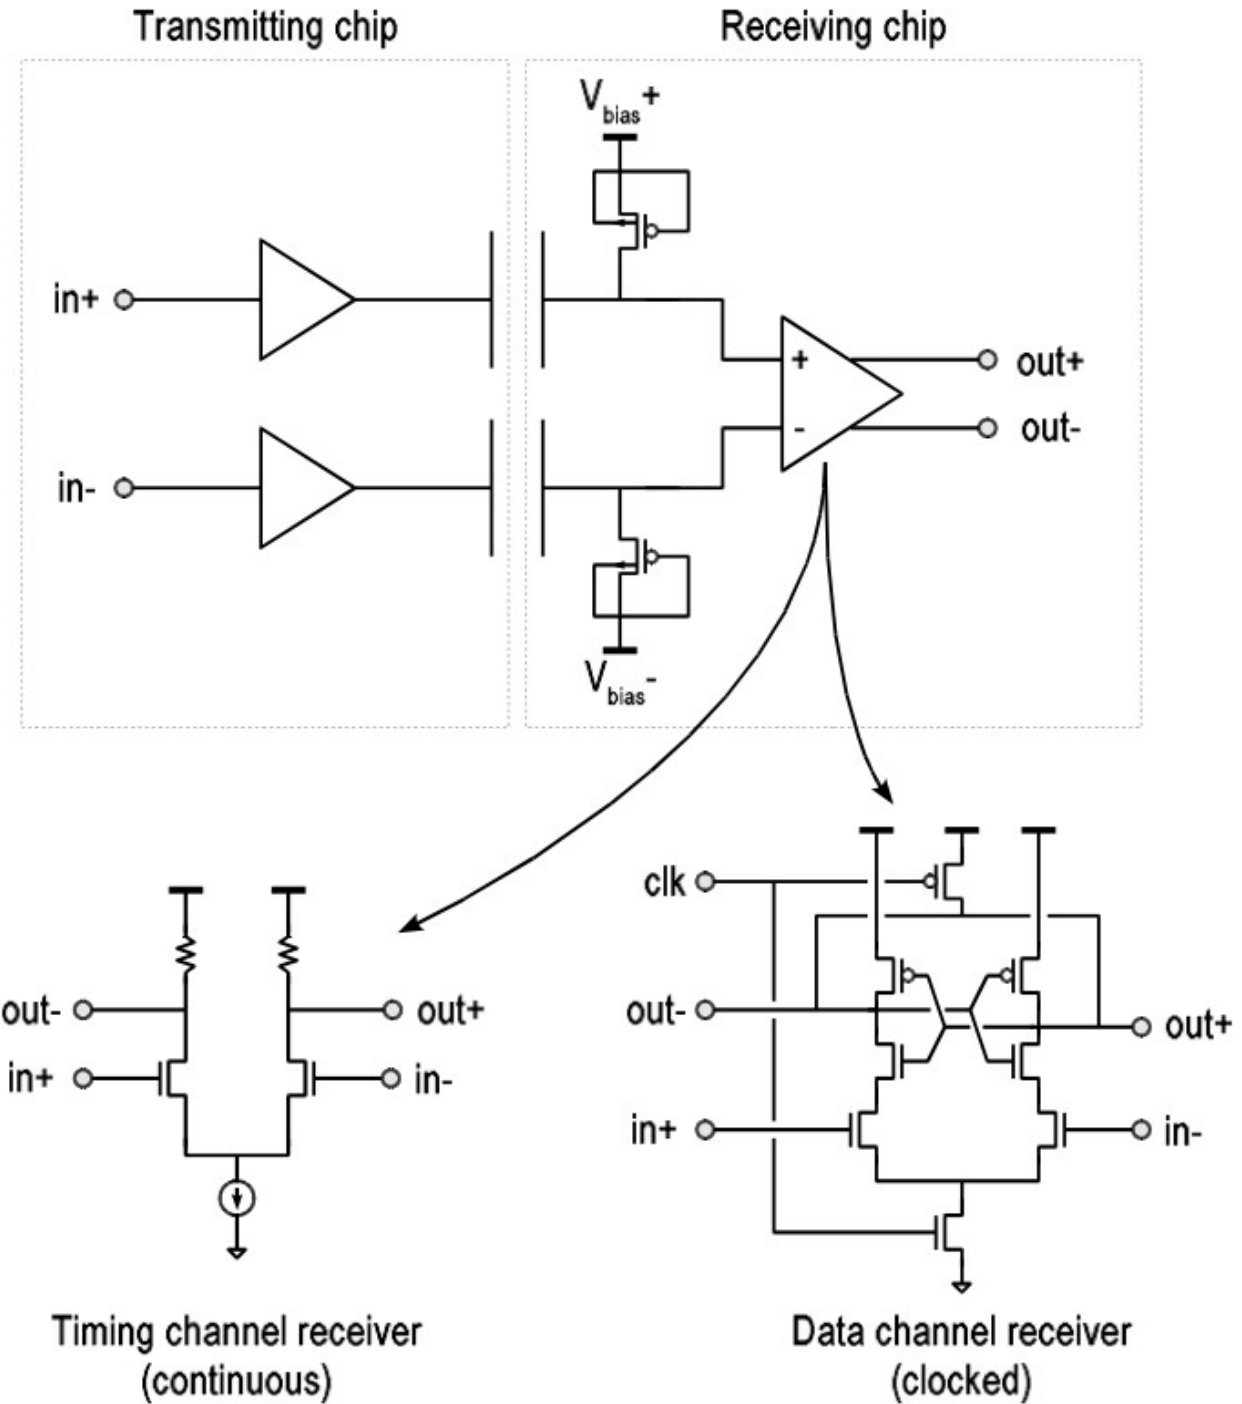
# Proximity Communication



- Avoids Off-Chip Wires
- Increases Bandwidth/Area
- Makes Chips Replaceable
- Enhances Testing Capability
- Enables Smaller Chips
- Obviates ESD Protection
- Shrinks Transceiver Circuits




# Simple Circuits:




# Proximity Packaging Challenges

- Performance is a function of  $Z$ ,  $\Psi$ ,  $\Phi$  misalignments




- With reasonable misalignment control tens of Tbps bandwidth per chip can be realized

# Alignment is Multi-Dimensional




# Alignment is the major challenge


- Must align chips in  $X$ ,  $Y$ ,  $\Theta$ ,  $Z$ ,  $\Psi$ ,  $\Phi$
- $X$ ,  $Y$ ,  $\Theta$  misalignments are corrected electronically



Measure...




...and correct...



...on-chip

# Steering Circuit One Receiver Pad Pitch

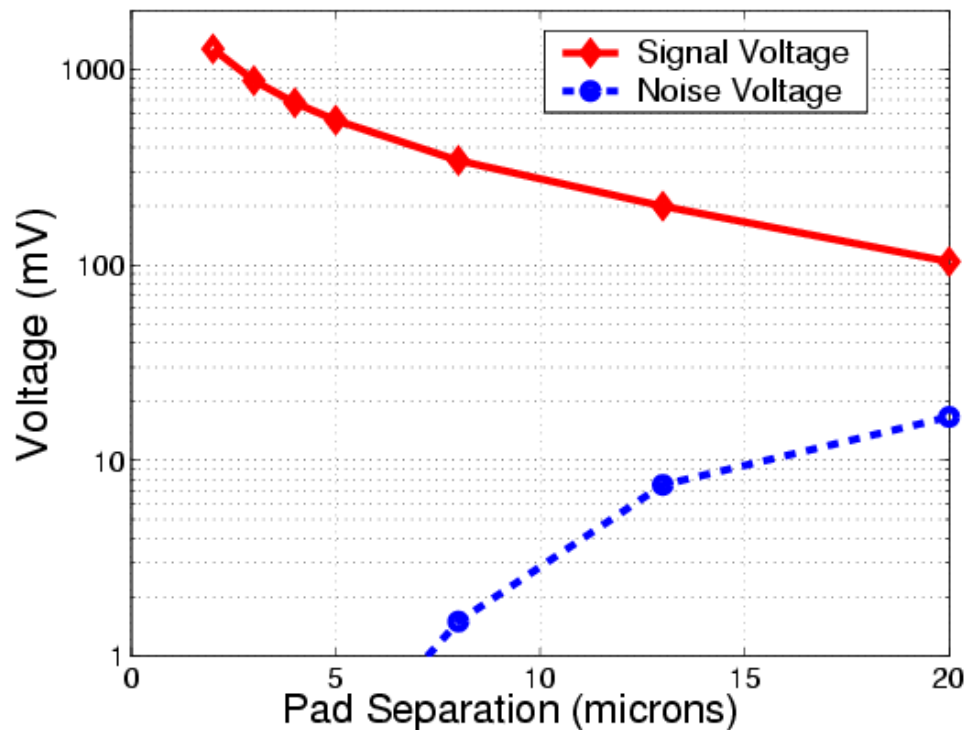


**Steering  
in two  
dimensions**

# Pads Cross-Section



# Signal and Noise Simulated Coupling




- Combining estimates for
- Channel speed
  - Receiver sensitivity
  - Signal vs. noise for pads
  - Clocking and overhead

We can estimate...

$$BandwidthDensity = \frac{77}{G^2} \frac{Tb/s}{mm^2}$$

Where G=pad separation, or gap in microns



# A tileable PxC block



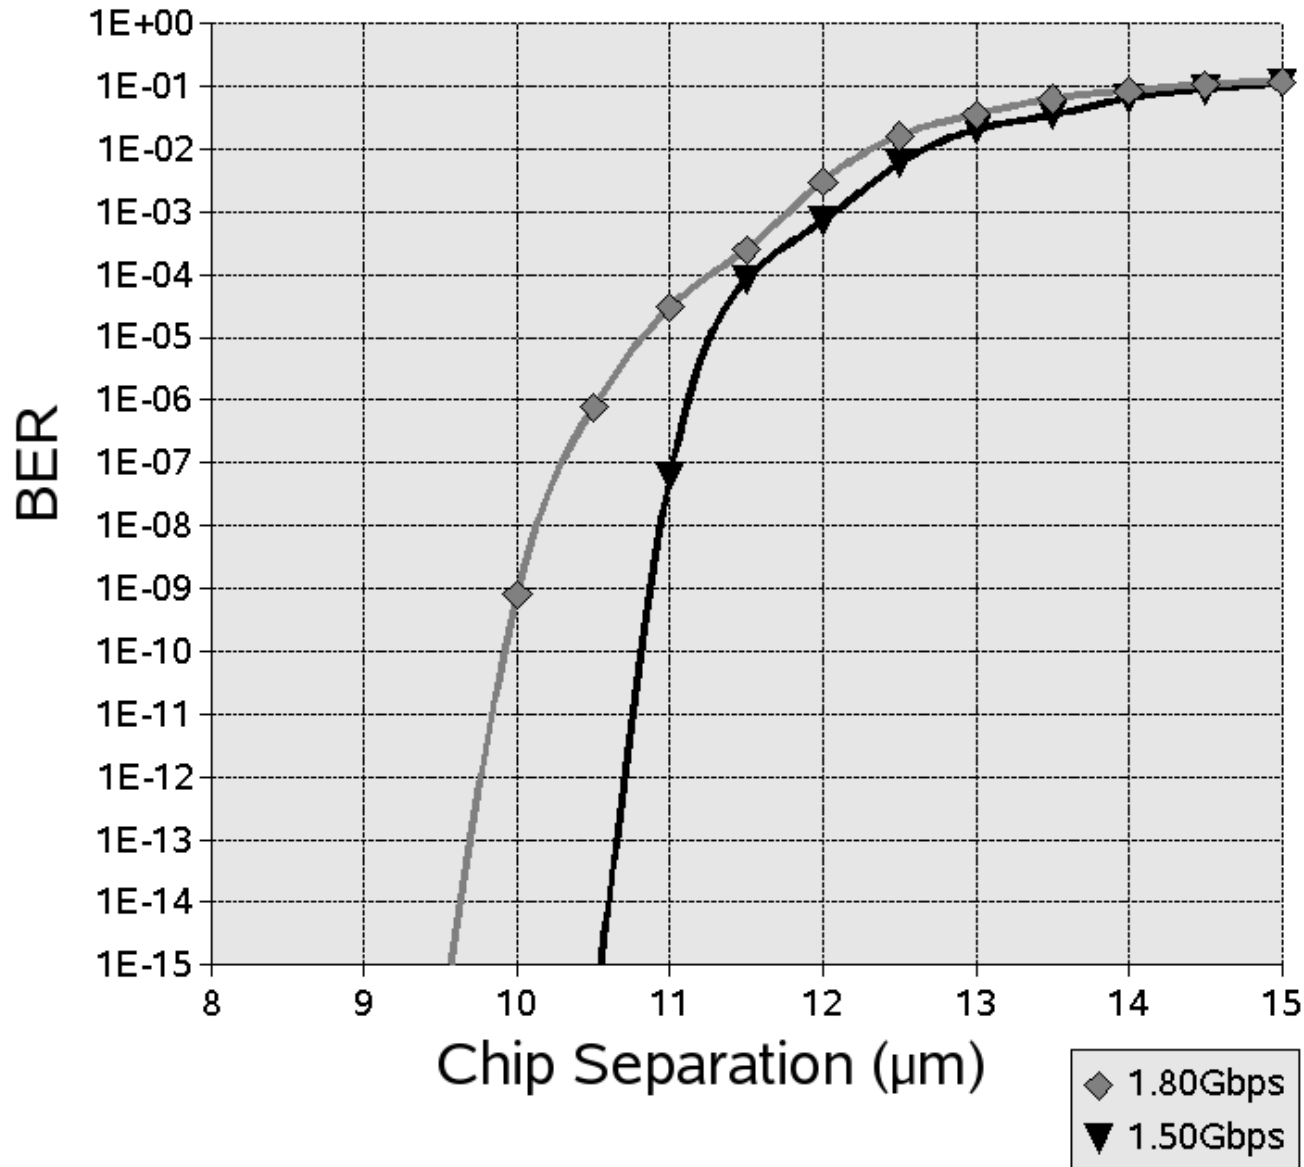
# Measured results

- TSMC 180nm CMOS
- 72 transmit, 72 receive channels
- 1.8 Gb/s per channel,  $10^{-15}$  bit error rate
- Aggregate 260Gb/s/chip, density 430Gb/s/mm<sup>2</sup>
- 3pJ/bit

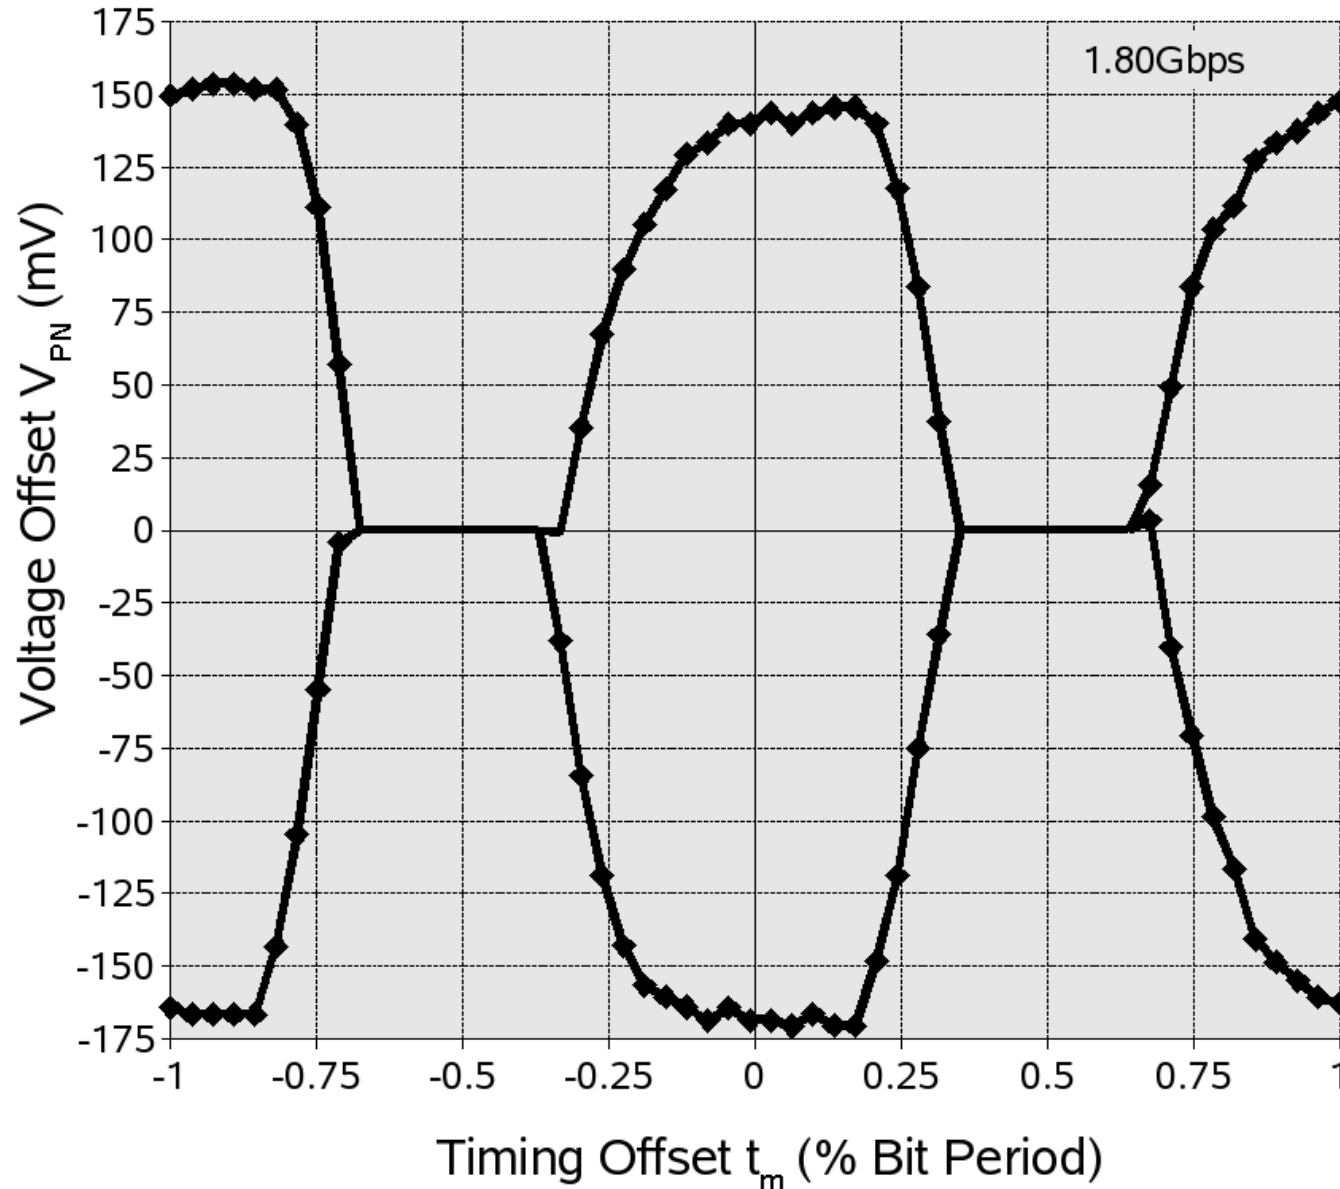
# Experimental Setup



# BER vs. chip separation



# Eye opening at 1.8Gb/s



# How do we multiply Moore's Law?

# The Key Idea in Moore's Law

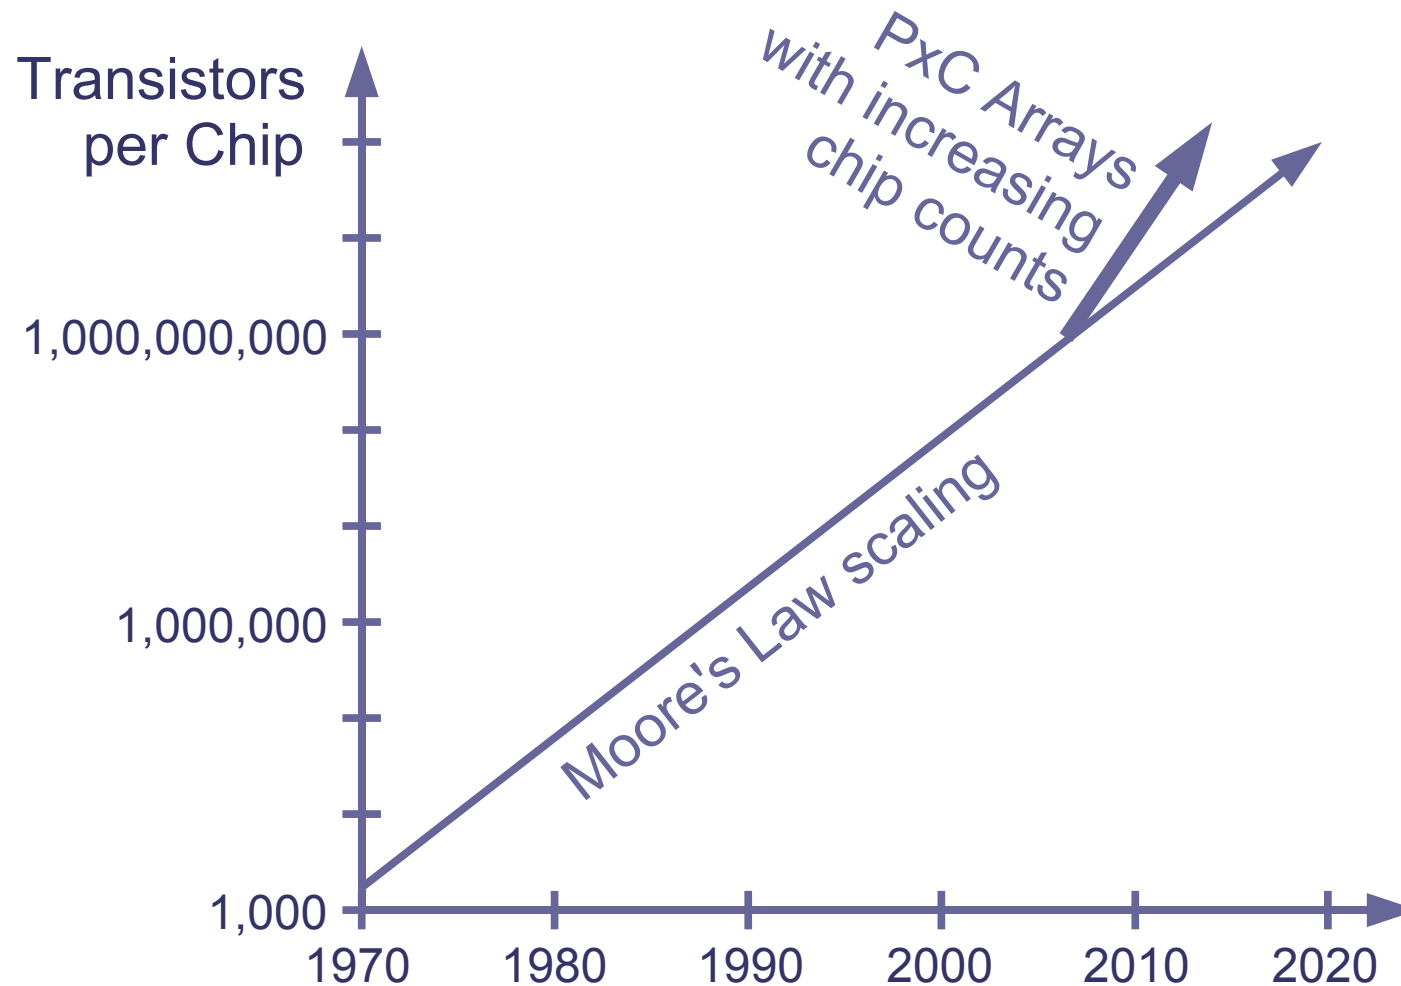
- Double number of transistors/chip (for same cost) every 24 months
  - > The principal driving force behind the past 40 years of integrated circuit industry advancement
  - > An amazing prediction in 1965 based on fewer than a hundred transistors/chip

# The Key Idea in Proximity Comm.

- We connect chips with enough bandwidth that they can perform as a single integrated chip
- Hence, PxC increases the effective number of transistors/chip over and above Moore's Law

# Multiplying Moore's Law

- Assuming Moore's Law continues




# What if Moore's Law stalls?

- Many have (incorrectly) predicted demise of Moore's Law
- Technical causes
  - > Short channel effects in transistors leading to too much leakage and hence power consumption
  - > Wire delay limiting performance
- Financial causes
  - > Fabs cost too much to yield a return on investment
    - 65nm fabs cost \$3 Billion to build (and going up 2x per generation)
  - > Chips cost too much to yield a return on investment

# Multiplying a stalled Moore's Law

- Proximity Communication keeps increasing transistors/chip without a fabrication contribution



# Summary

- Need for off-chip bandwidth motivates PxC
- Good mechanical alignment enables PxC and its tremendous bandwidth increase
- PxC multiplies Moore's Law by providing enough bandwidth to realize wafer-scale integration



# Multiplying Moore's Law with Proximity Communication

<http://research.sun.com/vlsi>



2007  
Sun Labs  
Open House



# References

- (1) D. Hopkins, et al., “Circuit Techniques to Enable 430Gb/s/mm<sup>2</sup> Proximity Communication,” *IEEE Int'l Solid-State Circuits Conference*, Feb. 2007.
- (2) R. Drost, et al., “Challenges in building a flat-bandwidth memory hierarchy for a large-scale computer with proximity communication,” *High Performance Interconnects, 2005. Proceedings. 13th Symposium on*, pp. 13-22, Aug. 2005.
- (3) Krste Asanović, et al., “The Landscape of Parallel Computing Research: A View from Berkeley,” *EECS Technical Report, in press*, December 3, 2006.
- (4) R. Drost, R. Ho, R. D. Hopkins, I. Sutherland, “Electronic Alignment for Proximity Communication,” *IEEE Int'l Solid-State Circuits Conference*, Feb. 2004.
- (5) R. Drost, R. D. Hopkins, I. Sutherland, “Proximity Communications,” *IEEE Custom Integrated Circuits Conference*, pp. 469-472, Sept. 2003.
- (6) J.L. Hennessy and D.A. Patterson, Computer Organization and Design, 2nd ed., Morgan Kaufmann Publishers, San Francisco, 1997.