



# Challenges of VLSI Interconnect

**Avinoam Kolodny**

**Technion – Israel Institute of Technology**



IBM HRL  
July 2009

# Outline

- ◆ **Interconnect as a basic system complexity problem**
- ◆ **Physical problem:**
  - The the poor scalability of metal wires**
- ◆ **Ways to address the problem:**
  - **Physical design of circuits and layout**
  - **System architecture directions**



# Technion's VLSI architecture research team:

## ◆ PIs

- *Israel Cidon (Networking)*
- *Ran Ginosar (VLSI)*
- *Idit Keidar (Dist. Systems)*
- *Isaac Kesslassy (Networking)*
- *Avinoam Kolodny (VLSI)*
- *Uri Weiser (Architecture)*

## ◆ Collaborators

- *Ronny Ronen*
- *Avi Mendelson*
- *David Goren*
- *Israel Wagner*
- *Eby Friedman*
- *Shmuel Wimer*

## Students:

*Evgeny Bolotin*  
*Reuven Dobkin*  
*Zvika Guz*  
*Ran Manevich*  
*Kostya Moiseev*  
*Arkadiy Morgenshtein*  
*Zigi Walter*  
*Anastasia Barger*  
*Shay Michaely*  
*Nir Magen*  
*Michael Moreinis*






# Connectivity and Complexity

## Challenge of System Complexity






# Interconnect : The old hidden problem





Intel 8088 processor  
**Single metal layer**

Perceived Solution:  
Systolic arrays



# The impact of a second metal layer






# Adding more metal layers...

## ◆ Wires dominate:

- Delay
- Power
- Noise
- Reliability
- Cost




BACK END

Cu  
interconnect

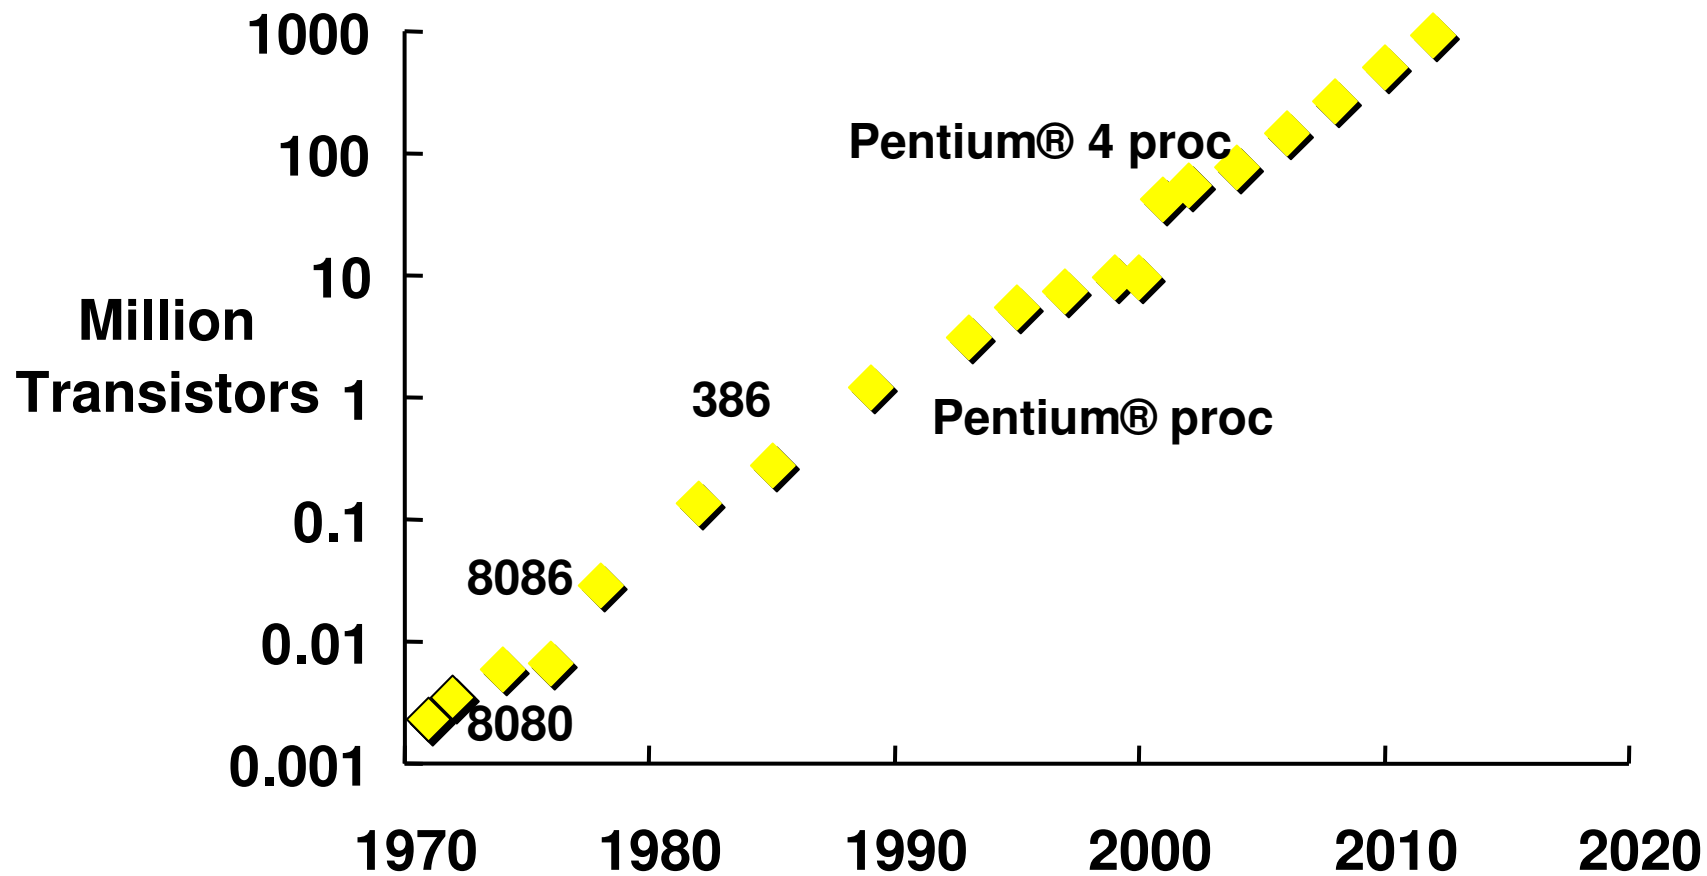
Low K  
dielectric

$\text{Si}_3\text{N}_4$   
Etch stop

FRONT END



# Moore's Law: Exponential growth in complexity by technology scaling





# Principles for dealing with complexity

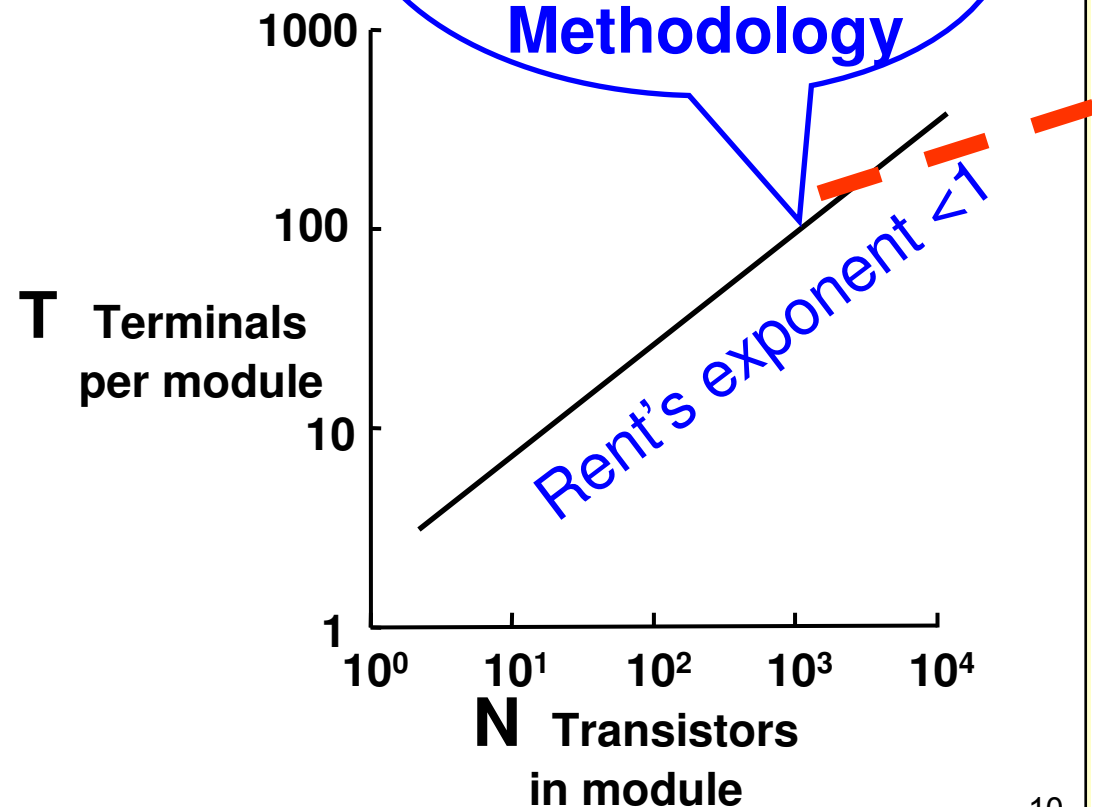
- **Abstraction**
- **Hierarchy**
- **Regularity**
- **Design Methodology**



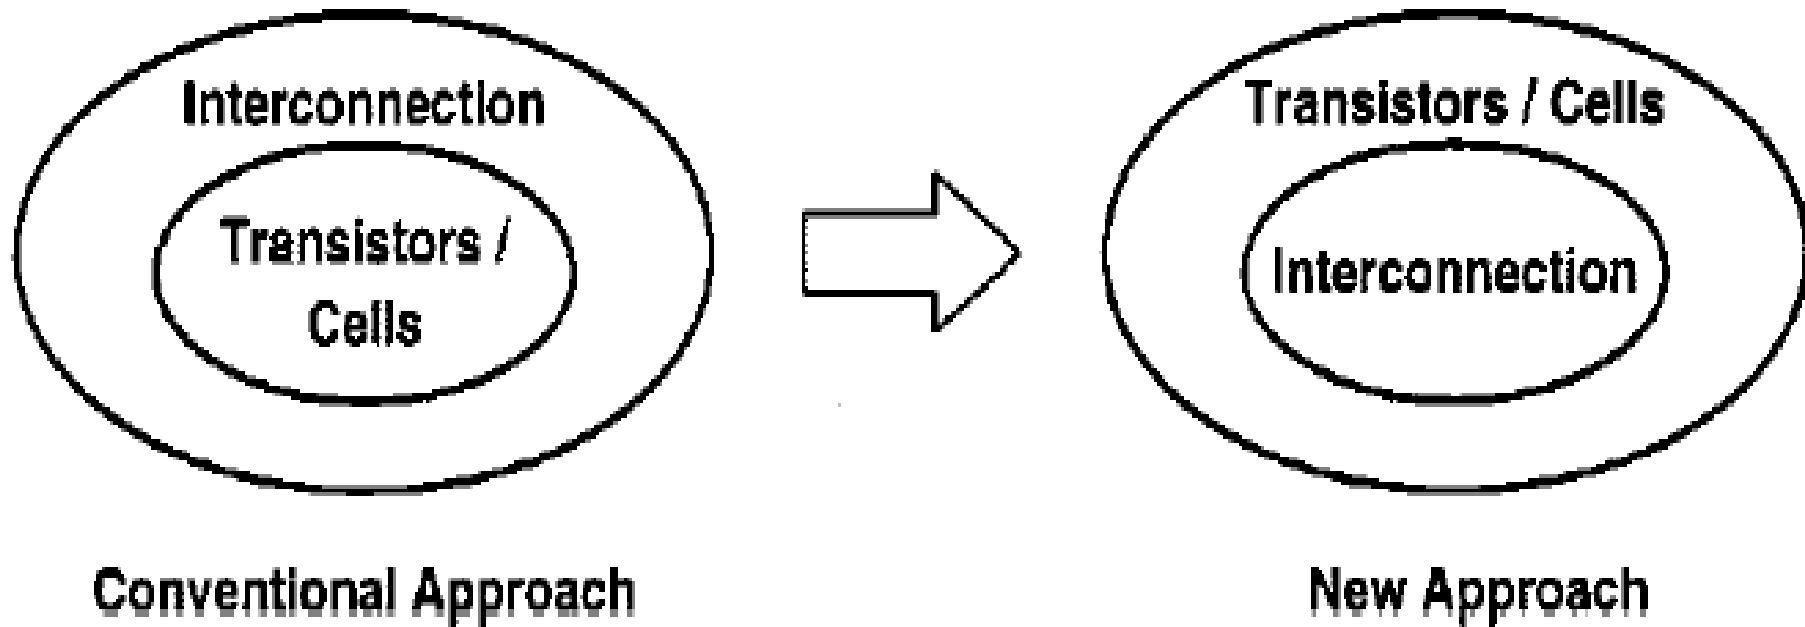
# Keeping up with Moore's Law: Rent's Rule

$$T = kN^r$$

- Abstraction
- Hierarchy
- Regularity
- Design Methodology



# The future of design flow?



- ◆ J. Cong : “Interconnect-centric design” Proc. IEEE 2001






# Physics of Technology Scaling


## ◆ For transistors, smaller is better!

- Speed *is proportional to  $L^{-1}$  or  $L^{-2}$*
- Power
- Lower cost, higher yield
- Increased system integration and reliability




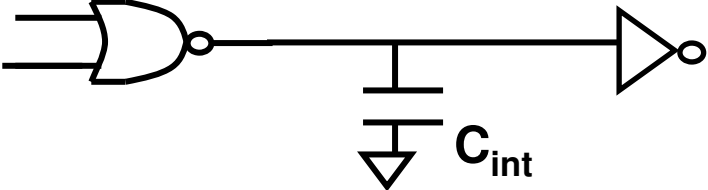
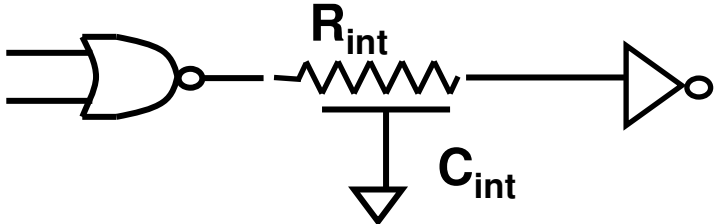
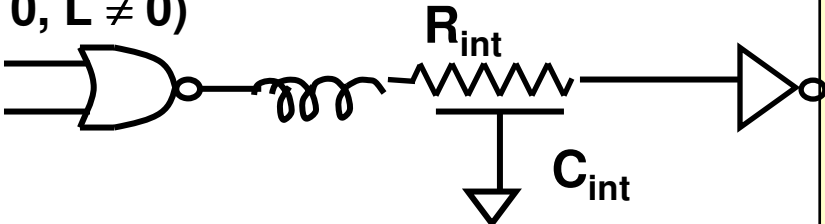
## ◆ For wires, smaller is worse...

*wire resistance grows*



# Evolution of interconnect models

Technology Evolution

- 1) "Ideal" Interconnect ( $R=0, C=0, L=0$ ) 
- 2) Capacitive interconnect ( $C \neq 0$ ) 
- 3) Resistive interconnect ( $C \neq 0, R \neq 0$ ) 
- 4) Inductive interconnect ( $R \neq 0, C \neq 0, L \neq 0$ ) 

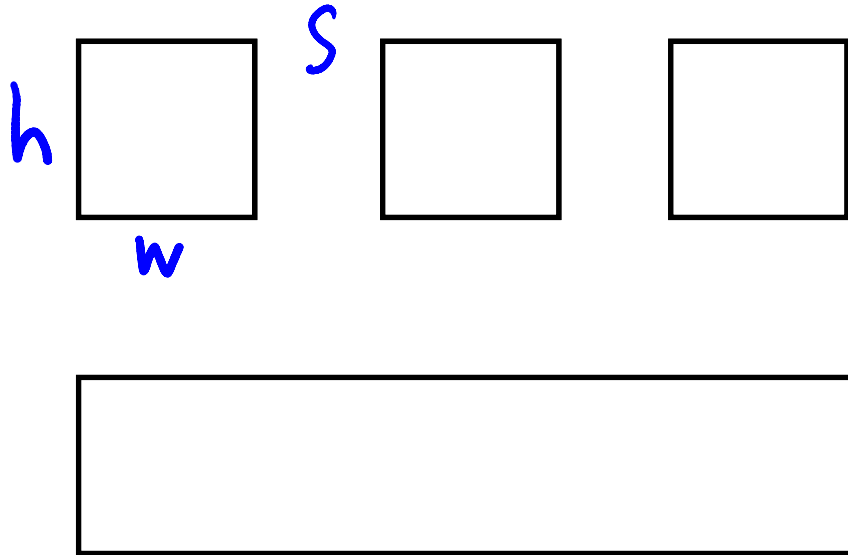


# Challenge of Interconnect Delay





# Nonuniform Scaling of Wires





The idea:

Shrink lateral dimensions – save area


Keep vertical dimensions – to avoid very high resistance




# Nonuniform Scaling of Wires



# Interconnect Capacitance $C = \frac{\epsilon A}{d}$




$\perp C_{\text{ground}}$



$\perp C_{\text{ground}}$

"Old" wires  
capacitance to ground was dominant



$C_{\text{cross}}$


"New" wires  
cross-capacitance is dominant






# “interconnect scaling – the real limiter”

Bohr, IEDM 95



$\tau = R_{\text{wire}} C_{\text{wire}}$   
 (time constant of the wire)  $\approx$  delay

$$\tau \approx \left( \rho \frac{L}{0.5P \cdot T} \right) \left( 2\epsilon_r \epsilon_0 \left( \frac{L \cdot 0.5P}{T} + \frac{L \cdot T}{0.5P} \right) \right) = 2\rho\epsilon_r\epsilon_0 L^2 \left( \frac{1}{T^2} + \frac{4}{P^2} \right)$$



# “interconnect scaling – the real limiter”

Bohr, IEDM 95; ITRS 1997




Figure 3 Calculated Gate and Interconnect Delay versus Technology Generation

Calculated gate and interconnect delay versus technology generation illustrating the dominance of interconnect delay over gate delay for aluminum metallization and silicon dioxide dielectrics as feature sizes approach 100 nm. Also shown is the decrease in interconnect delay and improved overall performance expected for copper and low κ dielectric constant insulators. <sup>1</sup>


# Local wires and Global wires

- ◆ **Local wire:**
  - Shrinks in length just like everything else
  - While transistors become faster, local wire delay remains unchanged (by simple scaling theory)
  
- ◆ **Global wire:**
  - Goes across the whole chip – does not scale!
  - Reflects new complexity added to the system!



# Bakoglu's solution: Repeaters

Bakoglu's classical derivation ED-32, 1985



# The global wire scaling problem

Gate delay gets better, wire delay gets worse

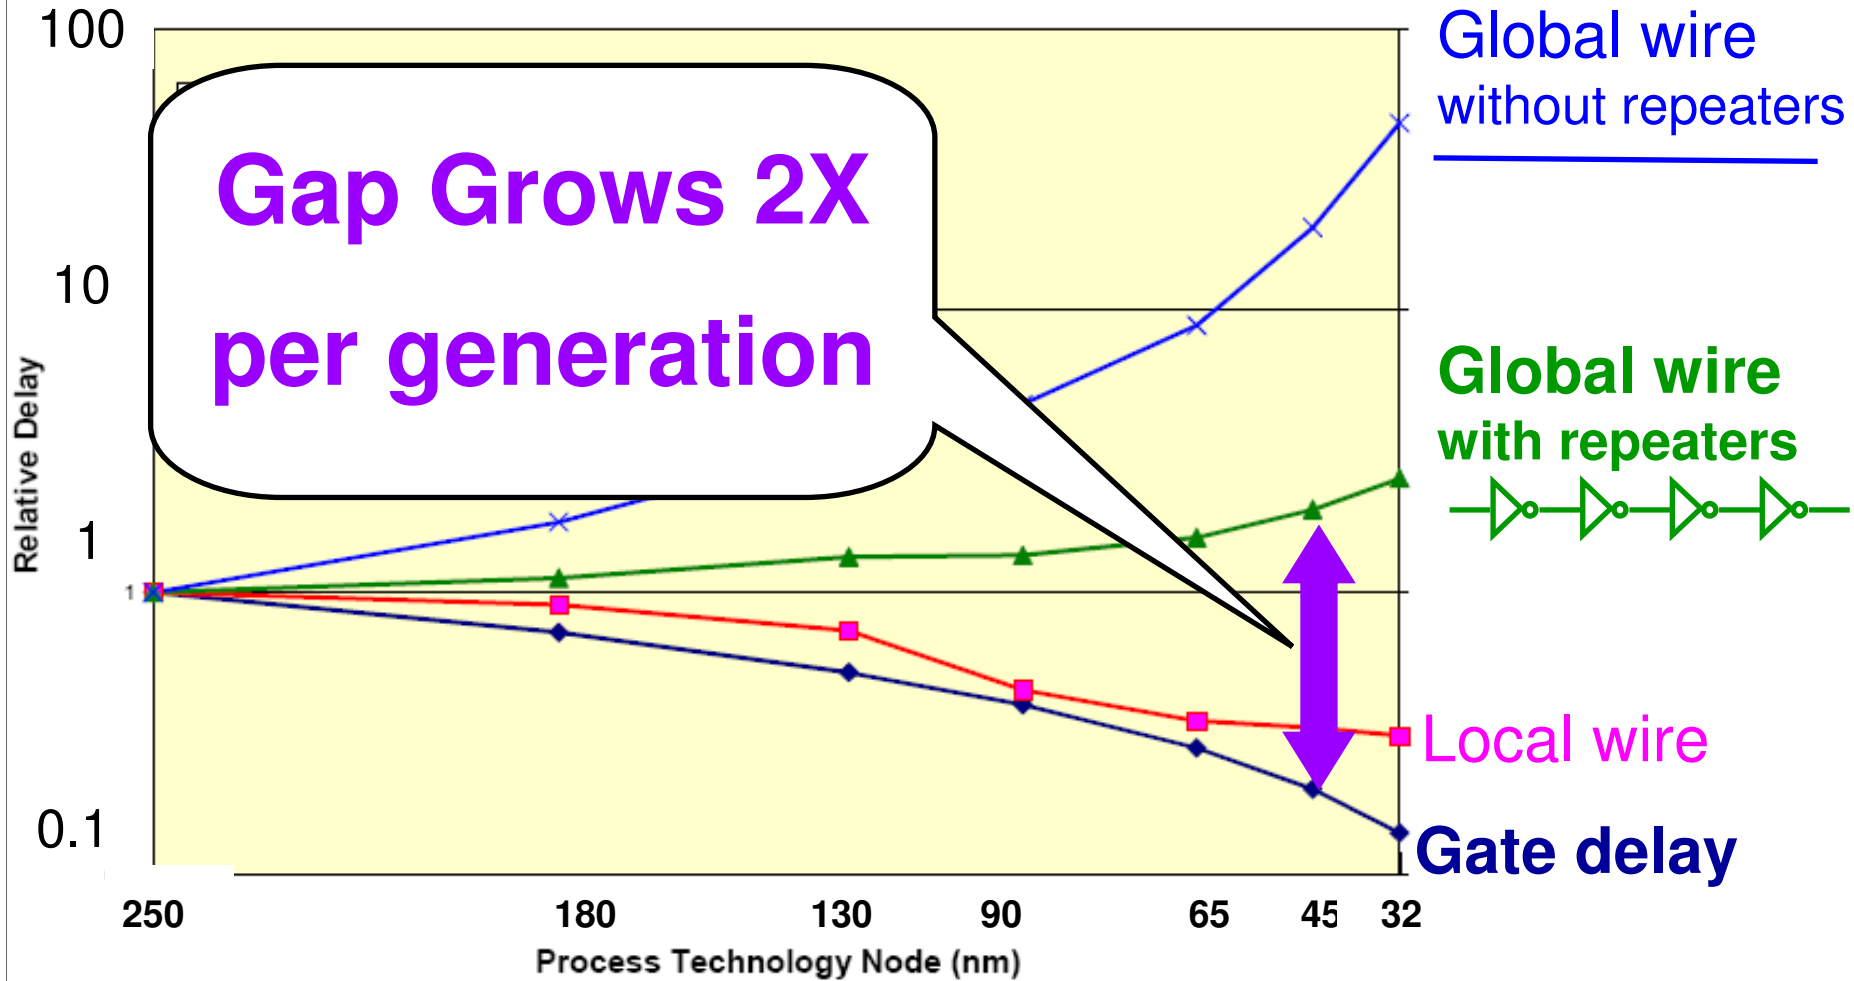



Figure 54 Delay for Metal 1 and Global Wiring versus Feature Size





# Delay of global wire is longer than a clock cycle

- ◆ Time for signal propagation across the die:
  - Today: 4-5 cycles?
  - In 10 years: >20 cycles?




**Fraction of chip reachable in 1 clock cycle**

Source: Keckler et al. ISSCC 2003




# Inverse Scaling: “fat wires”

- ◆ Thick & wide wires at the top metal layers:
  - ◆ Large cross section - Low R
  - ◆ Large spaces – Low C



# Inverse Scaling: “fat wires”

- ◆ Thick & wide wires at the top metal layers:
  - ◆ Large cross section - Low R
  - ◆ Large spaces – Low C

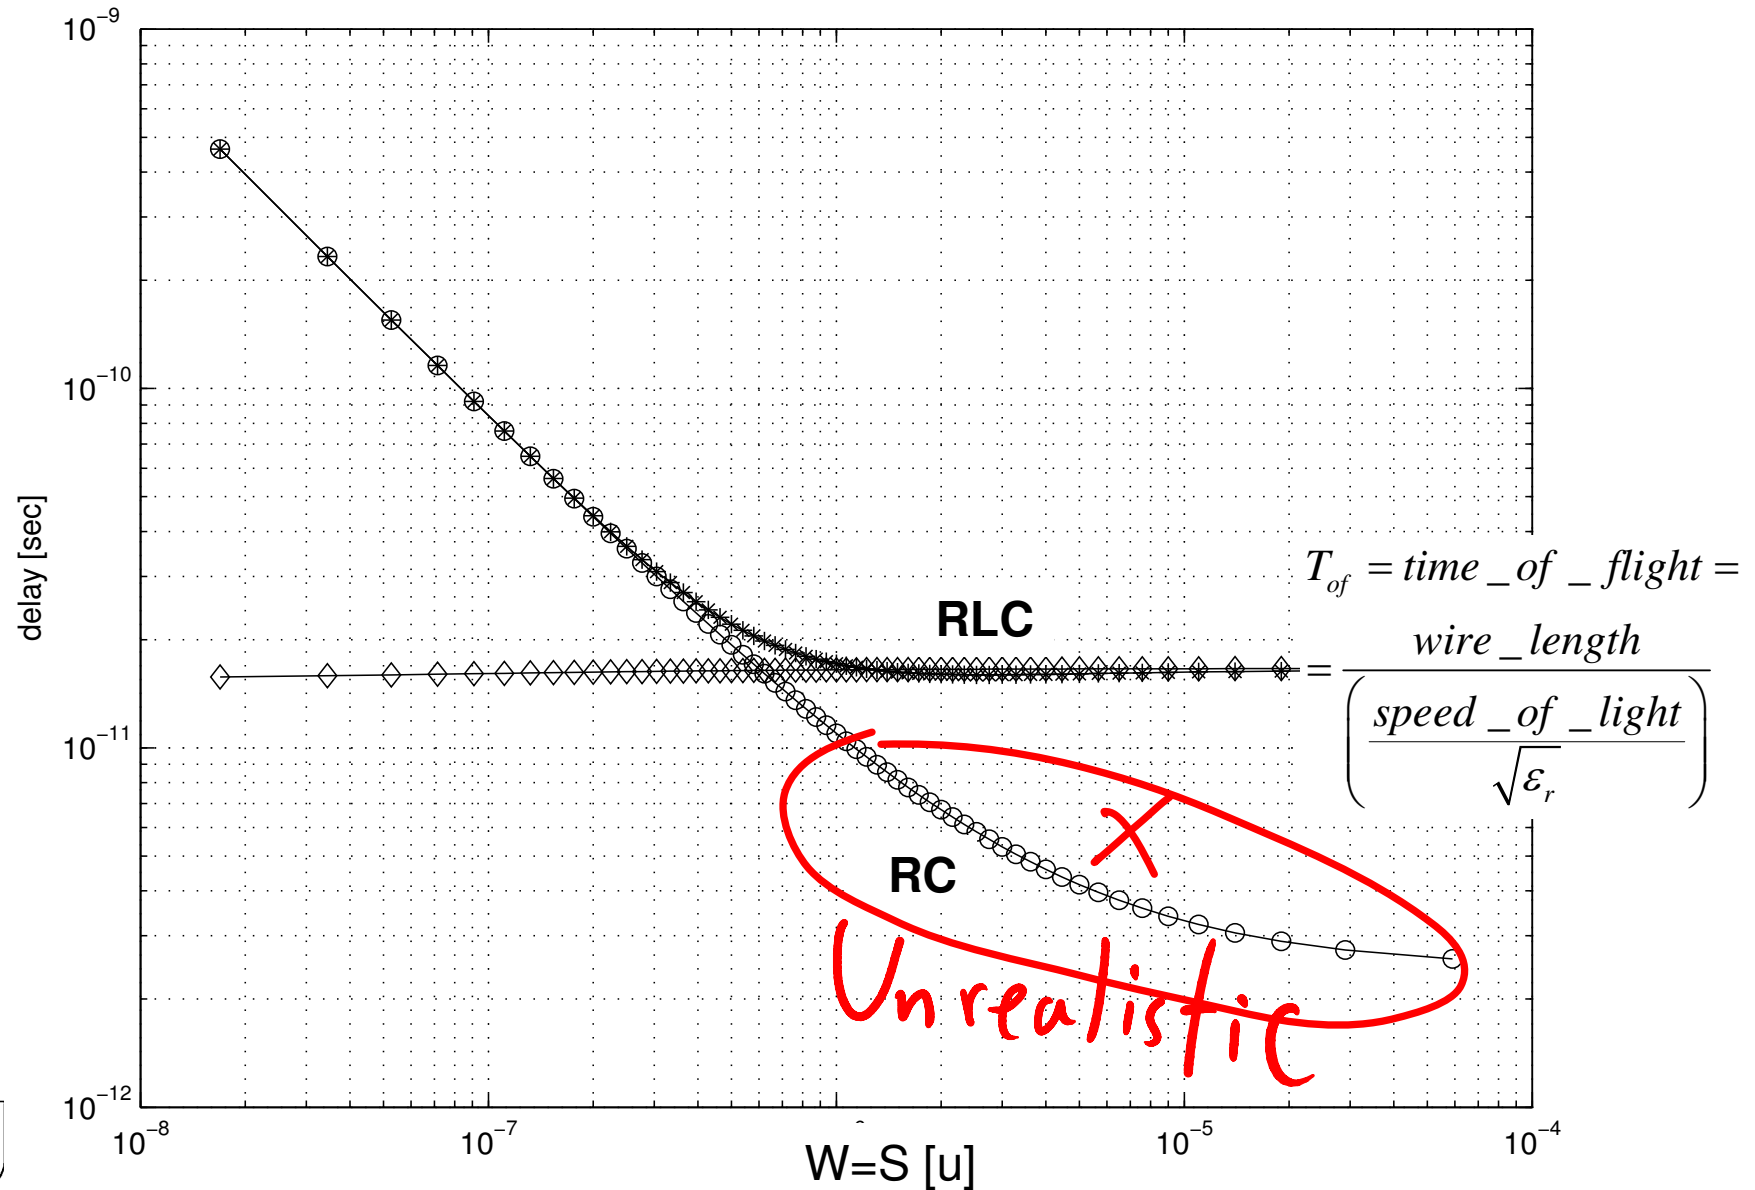


*Can we make  
fat wires  
as fast as we want?*



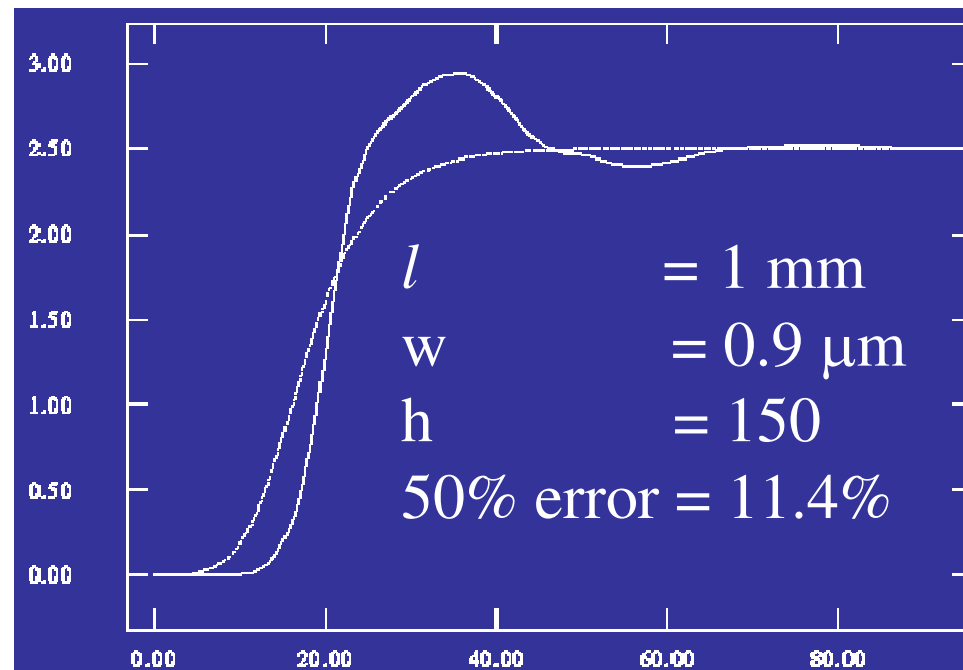

# What about the speed of light?

◆ Assume S=W



# Speed optimization in RLC lines

(Ismail & Friedman, ISCAS 99)





# RLC delay model characteristics

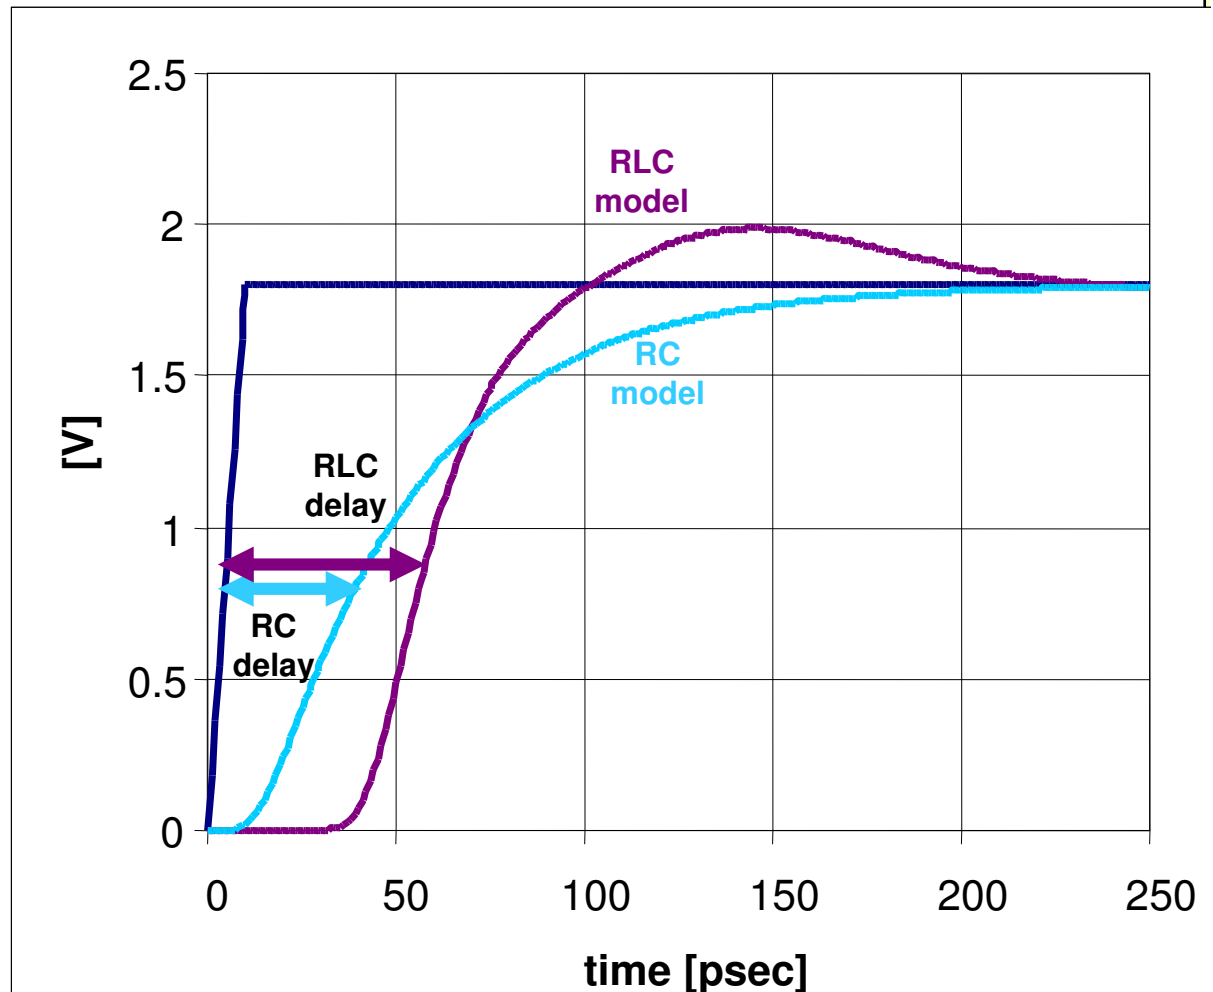
$$delay = \sqrt{LC} \left( e^{-2.9(\alpha_{asym} l)^{1.35}} l + 0.74 \alpha_{asym} l^2 \right)$$

$$\alpha_{asym} \triangleq \frac{R}{2\sqrt{\frac{L}{C}}}$$

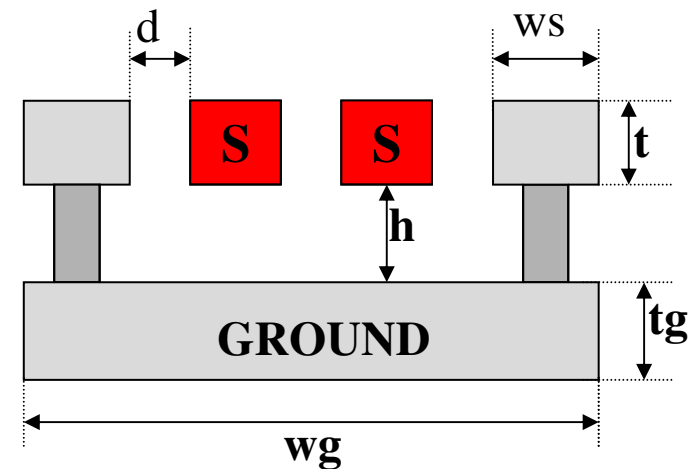


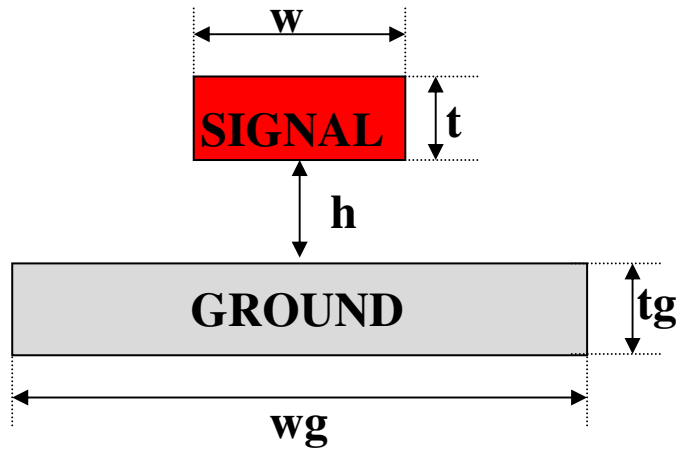
$L, C$  and  $R$  are per unit length  
 $l$  denotes wire length

- ◆ **Inductive effects:**
  - Longer delay
  - Steeper slope
  - overshoot

\* Eby G. Friedman, Yehea E. Ismail, On-chip inductance in high speed integrated circuits, 2001



# Fast wires must use transmission line layout

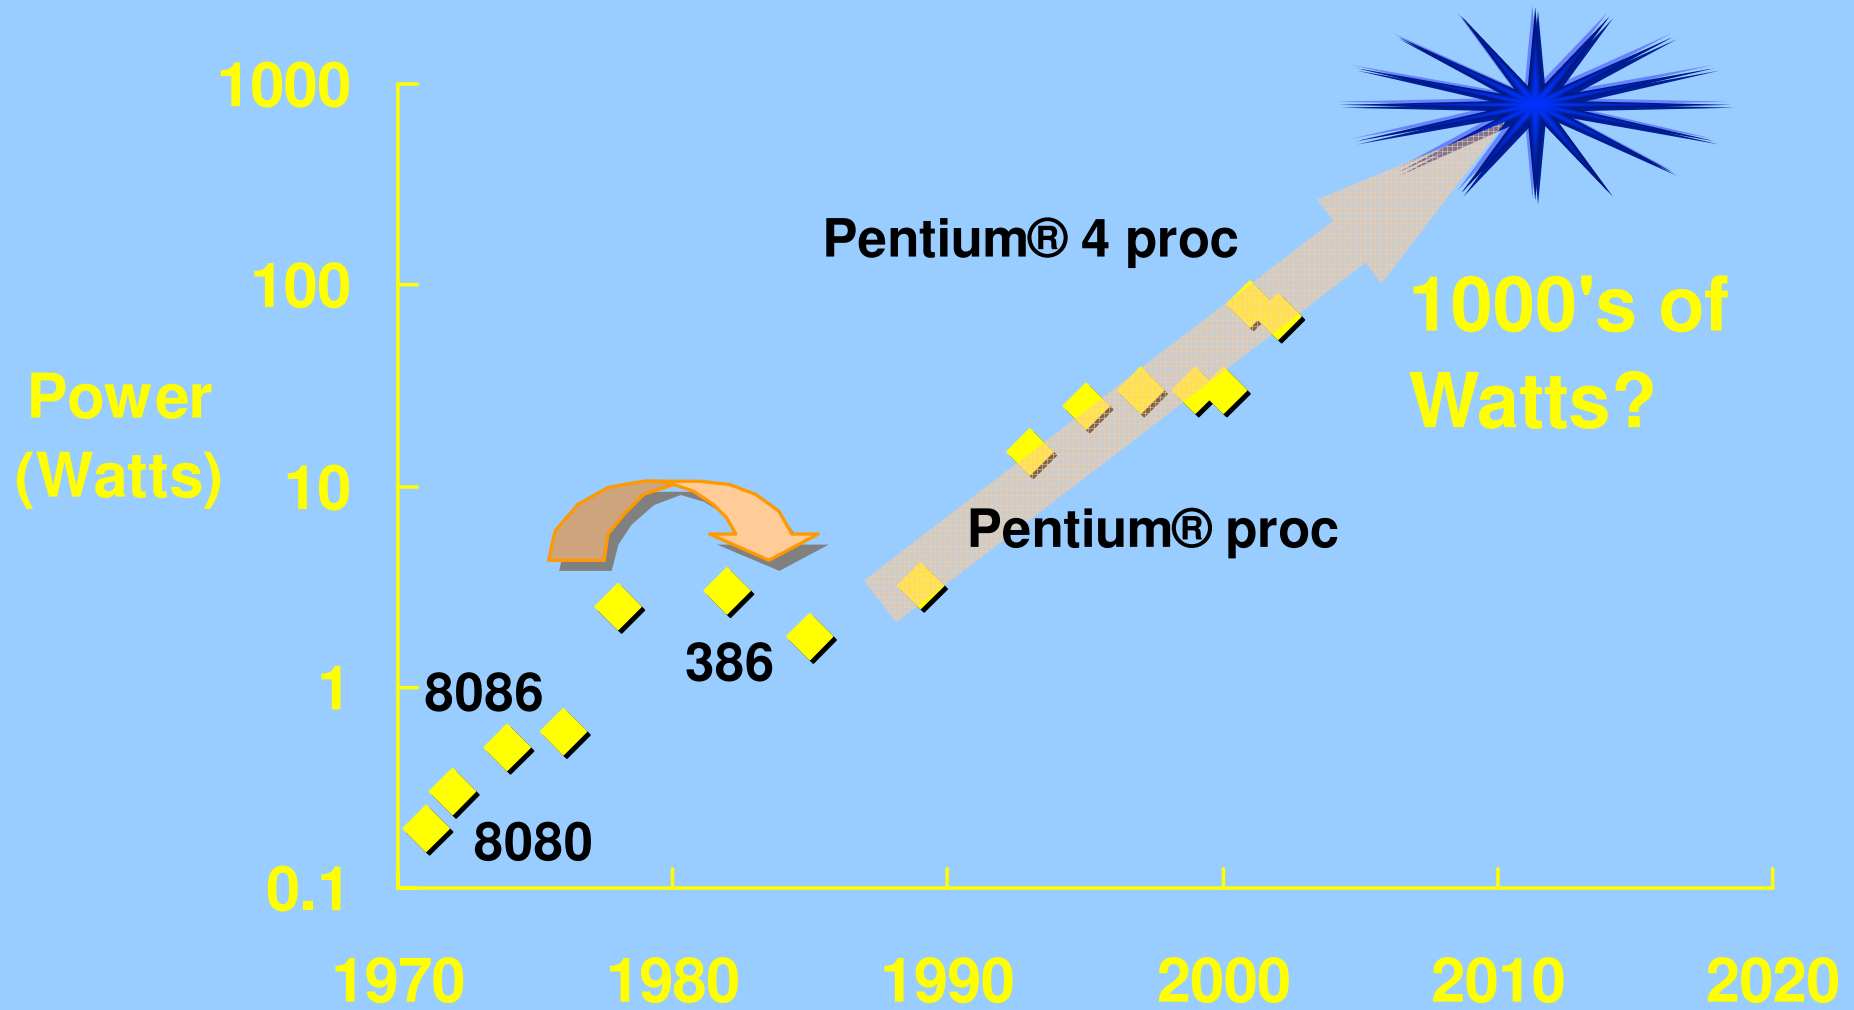


- ◆ Ground plane and/or wires provide *current return path*

# Challenge of Interconnect Power




# The infamous growth in processor power



# Interconnect power

- ◆ Definition: **Interconnect-Power**

Dynamic power consumption due to interconnect capacitance switching




- ◆ Dynamic power  $P = \sum A F_i \cdot C_i \cdot V^2 \cdot f$






# Interconnect Power in the Banias chip


**Interconnect power**



**Total dynamic power**



# The future of interconnect power



**Interconnect power grows to 65% -80% within 5 years**

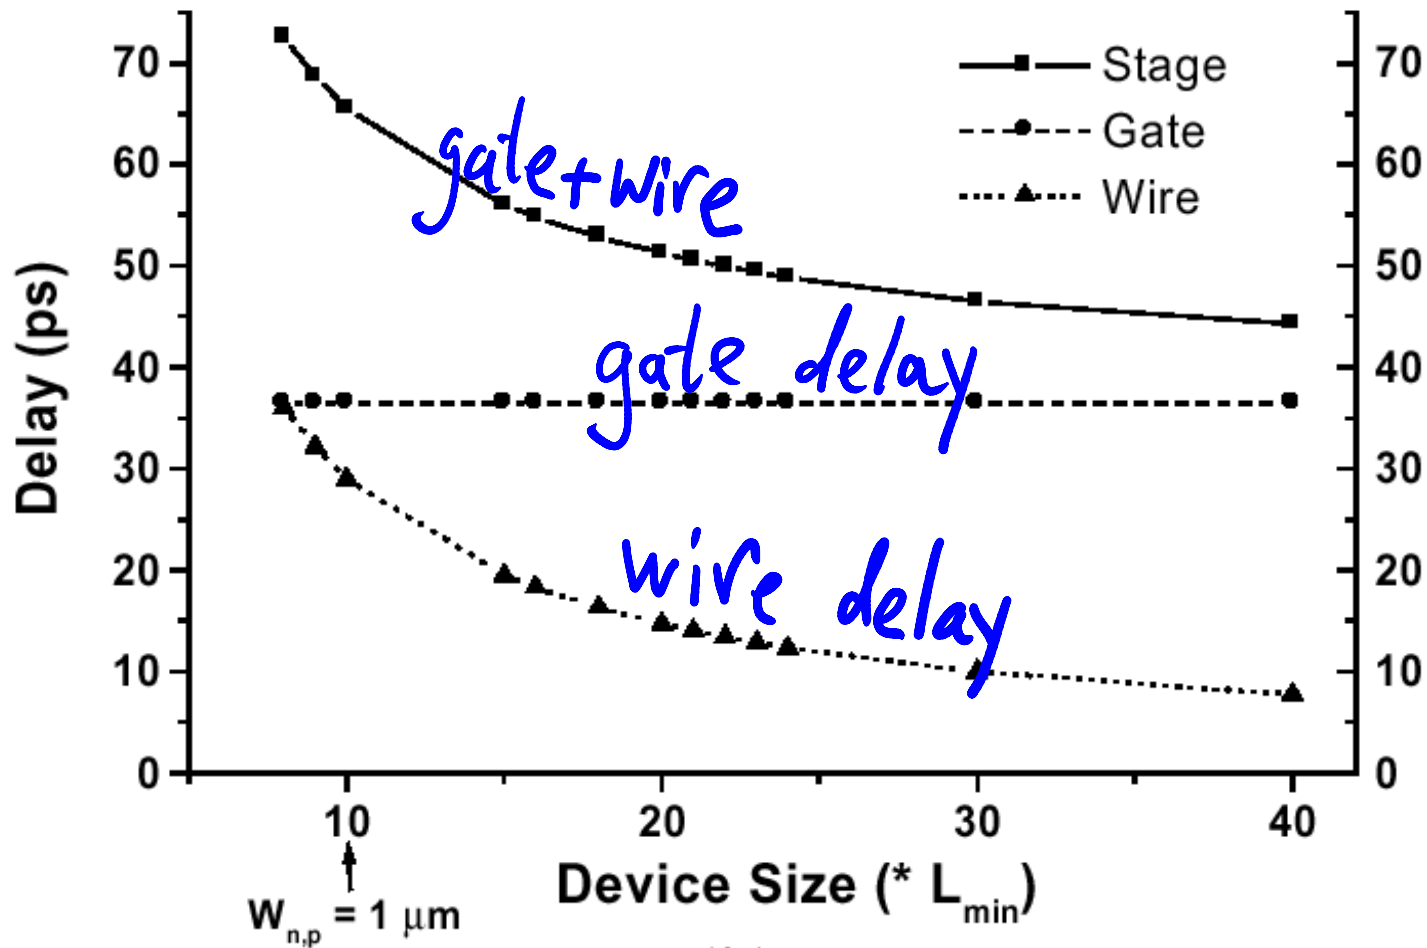

(using optimistic interconnect scaling assumptions for a uniprocessor)

**Global interconnect causes significant power dissipation**



# Wires can be blamed for even more power...

Because designers tend to  
oversize gates!





**Other Challenges:  
Interconnect Noise,  
Reliability,  
Cost**








# Design techniques for handling interconnect (some research contributions)



# Data Rate Optimization in an Interconnect Channel


$$\text{Data Rate} = N \cdot f = N \cdot \frac{1}{\text{delay}}$$



- ◆ Increase  $N$  by:
  - making the wires **narrow** (small  $W$ ),
  - and **dense** (small  $S$ )
- ◆ What will happen to the delay?



# Should all wires be the same? How about optimizing individual widths and spaces?



\* S. Wimer, S. Michaely, K. Moiseev and A. Kolodny,  
"Optimal Bus Sizing in Migration of Processor  
Design",  
IEEE Transactions on Circuits and Systems – I, vol.  
53, no. 5, May 2006.

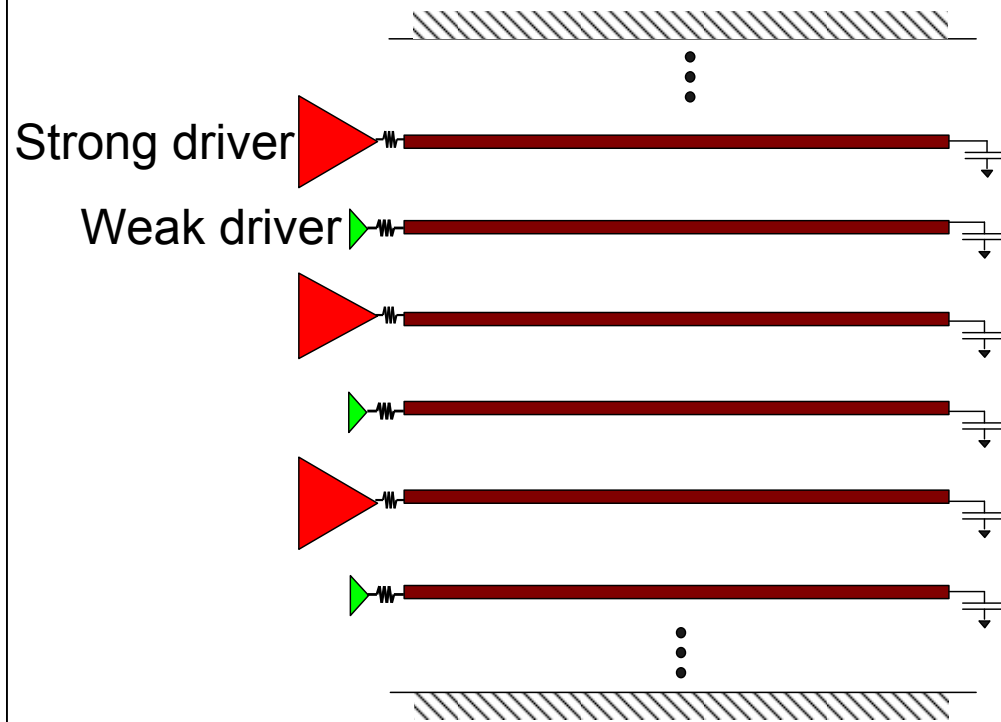
$$\sum_{j=1}^n W_j + \sum_{j=0}^n S_j = A$$

A is a fixed constraint






# Wire Reordering: Which order is better?



Worst order




Best order !



# Optimal order theorem


- ◆ given an interconnect channel whose wires are of uniform width  $W$ , 'Symmetric Hill' order of signals yields minimum total sum of delays (after spacing optimization).



\* K. Moiseev, S. Wimer and A. Kolodny, "Timing Optimization of Interconnect by Simultaneous Net-Ordering, Wire Sizing and Spacing," *INTEGRATION*, 2007.



# Wire spacing to reduce power?



**Average saving results: 14.3% for ASIC blocks <sup>1</sup>**

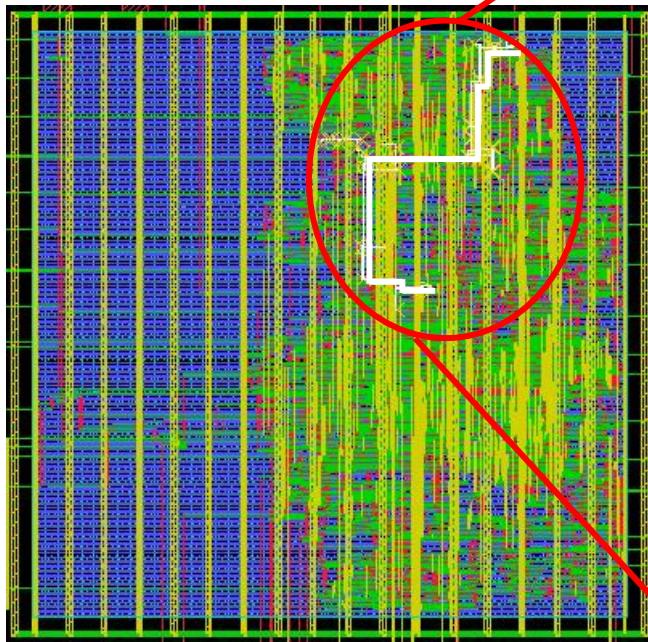
1 - Estimated based on clock interconnect power





\* N. Magen, A. Kolodny, U. Weiser and N. Shamir, "Interconnect-power dissipation in a Microprocessor," SLIP 2004.

# Logic with Wires

## Common Example



UART design





# The Interconnect Wall

Logic w/o wires

Long wires

Logic Gate Sizing

Interconnect Optimization



Logical Effort  
gate sizing

Repeater  
Insertion


$$g_i \cdot h_i = g_{i+1} \cdot h_{i+1}$$

$$x_{opt} = \frac{C_i}{C_0} = \sqrt{\frac{R_0 \cdot C_{w_i}}{R_{w_i} \cdot C_0}} \quad K = \sqrt{\frac{0.4 \cdot R_w \cdot C_w}{0.7 \cdot R_0 \cdot C_0}}$$



# Breaking The Wall: Logic Gates as Repeaters - LGR

*“Where should the gates be located (along the wire)?”*




\* M. Moreinis, A. Morgenshtein, I. Wagner, and A. Kolodny, “Logic Gates as Repeaters (LGR) for Area-Efficient Timing Optimization,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 14, no. 11, pp. 1276-1281, November 2006



# Breaking The Wall: Unified Logical Effort

Logic w/o wires

Long wires



## Challenges:

*Gate placements*

*Gate sizes*

*Number of gates, repeaters*

\* A. Morgenshtein et al., "Unified Logical effort – speed optimization of logic with interconnect", *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, in press.







# Some Radical Approaches:

**3D integration**

**CMPs** – chip multiprocessors


**Network on chip**

**Photonic interconnect**



# 3-D Integration

die side-length reduction  
4 planes ~50%




- Area =  $L^2$
- *Corner to corner distance =  $2L$*

- Same Area
- *Corner to corner distance  $\approx L$*



# Cross-section of a 3-D Integrated Circuit

- **Plane bonding**
  - Back to face
  - Face to face
- **Bonding materials**
  - Adhesive polymers
  - Metal pads (e.g., copper)
- **Bonding process involves**
  - Compression at elevated temperatures
  - Wafer thinning



\* R. J. Gutmann *et al.*, "Three-dimensional (3D) ICs: A Technology Platform for Integrated Systems and Opportunities for New Polymeric Adhesives," *Proceedings of the Conference on Polymers and Adhesives in Microelectronics and Photonics*, pp. 173-180, October 2001




# Multi-integration of 3 – D Systems-on-Chip


- **Integration of**
  - **Circuits from different fabrication processes**
  - **Non-silicon technologies**
  - **Non-electrical systems**




Image sensor array



Amplifier and ADC



Data latch and masking



Processor array and output circuit

• M. Koyanagi *et al.*, "Future System-on-Silicon LSI Chips,"  
*IEEE Micro*, Vol. 18, No. 4, pp. 17-22, July/August 1998



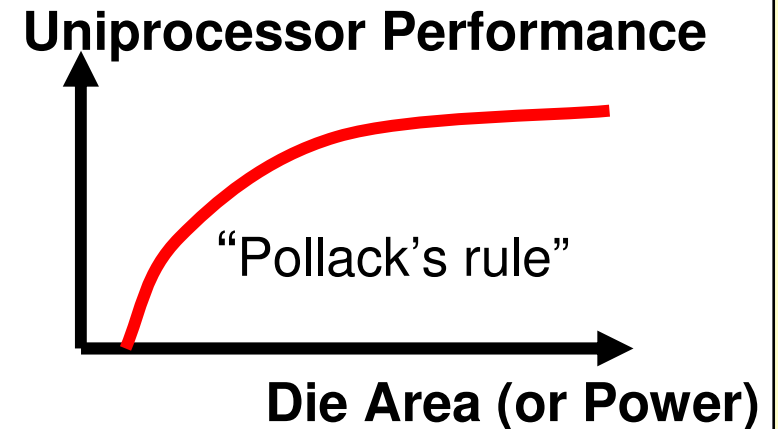
# Chip Multi-Processors

- ◆ **Uniprocessors cannot provide Power-efficient performance growth**

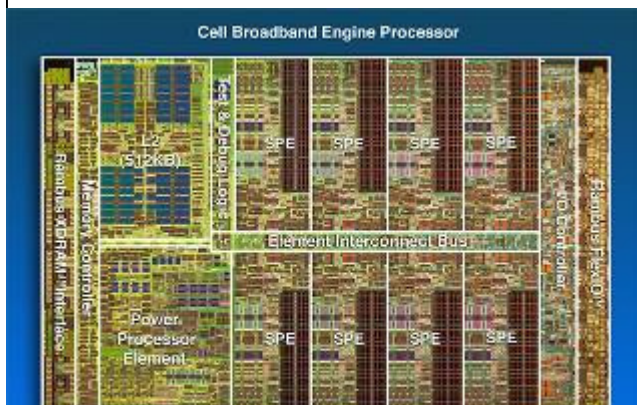
- Interconnect dominates dynamic power
- Global wire delay doesn't scale
- Instruction-level parallelism is limited

- ◆ **Power-efficiency requires many parallel local computations**

- Chip Multi Processors (CMP)
- Thread-Level Parallelism (TLP)




(F. Pollack. Micro 32, 1999)



# Future of VLSI architecture - CMP

(Dally 1999, Horowitz 2001)

- ◆ System requirement: Power-efficient performance growth
- ◆ Implications:
  - Chip Multi Processors (CMP)
  - Thread-Level Parallelism (TLP)
- ◆ Local memories
  - Memory is interconnection in time?
- ◆ Explicit communication






Possible future paradigm shift in VLSI:  
**Network on Chip (NoC)**

From: **Dedicated signal wires**



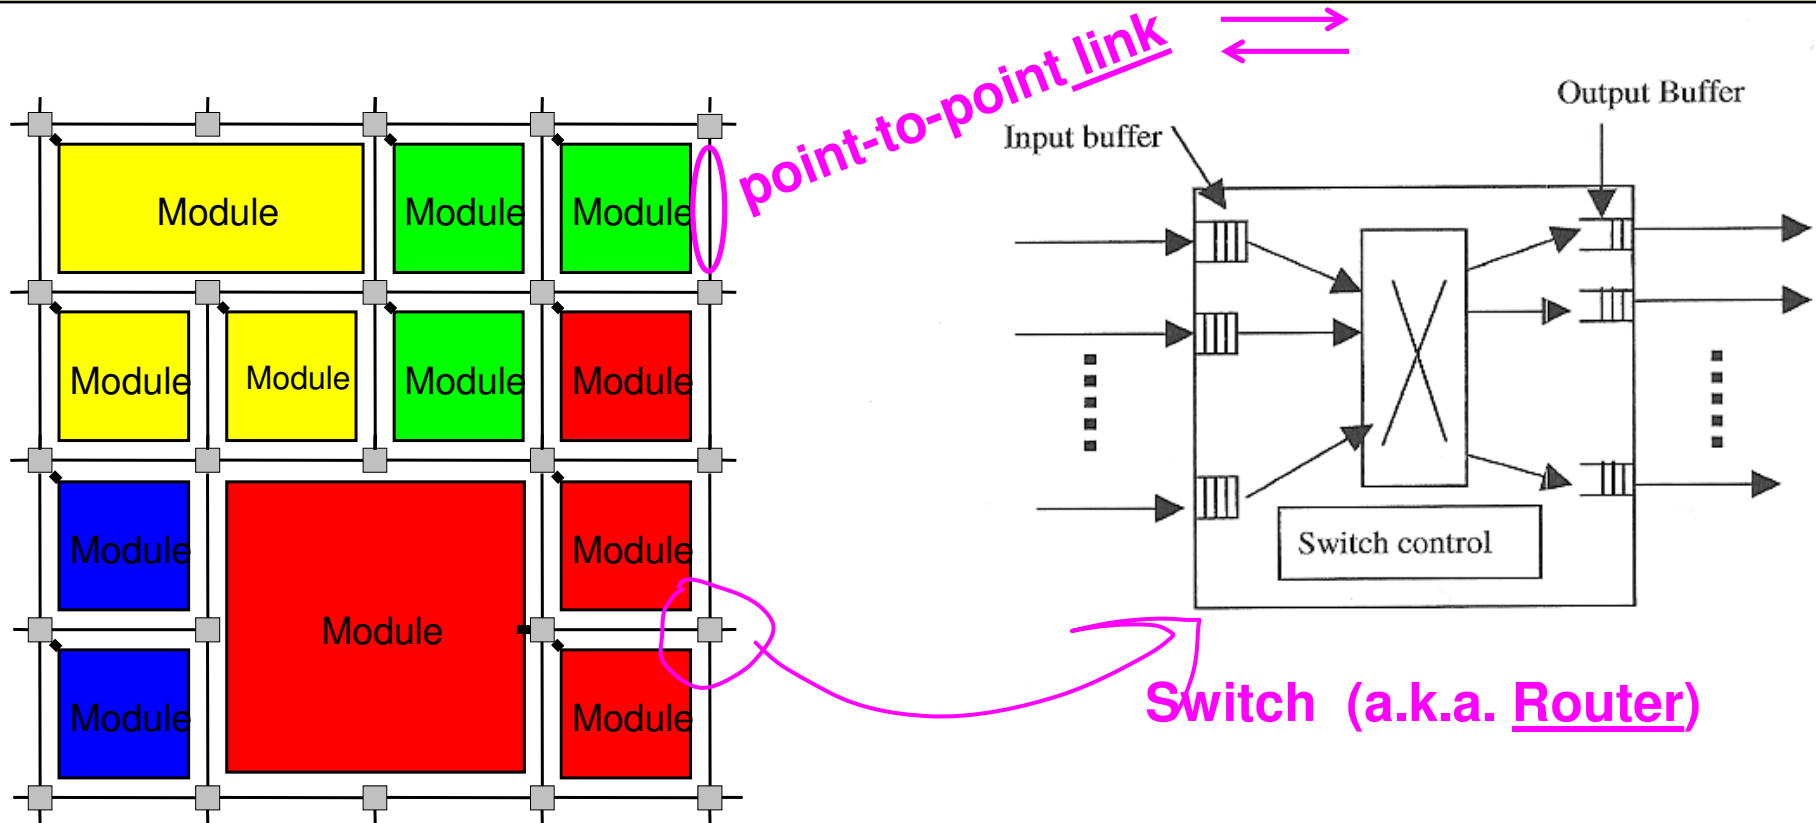
to: **Shared network**



□ Computing Resource    ■ Network switch    — Network link    53




# NoC Essentials



- ◆ **Communication by packets of bits**
- ◆ **Routing of packets through several hops, via switches**
- ◆ **Parallelism**
- ◆ **Efficient sharing of wires**



# Evolution or Paradigm Shift?




# Past Examples of Paradigm Shifts in VLSI

## The Microprocessor

From: **Hard-wired state machines**

To: **Programmable chips**

- ◆ Created a new computer industry




## Logic Synthesis

From: **Schematic entry**

To: **HDLs and Cell libraries**

- ◆ Logic designers became programmers
- ◆ Enabled ASIC industry and Fab-less companies
- ◆ “System-on-Chip”



*successful*

# Characteristics of a Paradigm Shift

- ◆ Solves a critical problem (or several problems)
- ◆ Step-up in abstraction
- ◆ Design is affected:
  - Design becomes more restricted
  - New tools
  - The changes enable higher complexity and capacity
  - Jump in design productivity
- ◆ Initially: skepticism. Finally: change of mindset!


*What are the problems addressed by NoC?*



# Critical Problems Addressed by NoC

1) Global interconnect design problem:  
delay, power, noise, scalability, reliability

2) System integration  
productivity problem



3) Chip Multi Processors  
(key to power-efficient computing)



# 1(a): NoC and Global Wire Delay


Long wire delay is dominated by Resistance

---


Add repeaters



Repeaters become latches (with clock frequency scaling)




Latches evolve to NoC routers




Source: W. Dally

# 1(b): Wire Design for NoC



- ◆ **NoC links:**
  - Regular
  - Point-to-point (no fanout tree)
  - Can use transmission-line layout
  - Well-defined current return path
- ◆ Can be optimized for noise / speed / power
  - Low swing, current mode, ....






## 2: NoC and the Engineering Productivity Problem

- ◆ **NoC eliminates ad-hoc global wire engineering**
- ◆ **NoC separates computation from communication**
  - NoC supports modularity and reuse of cores
- ◆ **NoC is a platform for system integration, debugging and testing**

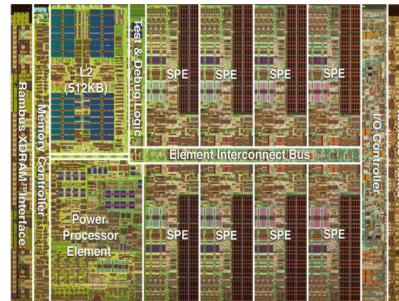


# 3: NoC and CMP


- ◆ Network is a natural choice for multiple cores!




**Montecito**  
Intel 2004




**CELL BE**  
IBM 2005



**Terascale**  
Intel Polaris 2007



**Niagara**  
Sun 2004




**Barcelona**  
AMD 2007




# Why Now is the Time for NoC?



Difficulty of DSM wire design




Productivity pressure



CMPs




# Combining 3-D and NoC?



- ◆ Each plane has a dedicated NoC
- or
- ◆ Dedicated NoC planes



# Hybrid Optical 3-D NoC?







## Summary:

- Interconnect is a major challenge in modern VLSI
- Exciting R&D opportunities!