

Digital Design for Low Power Systems

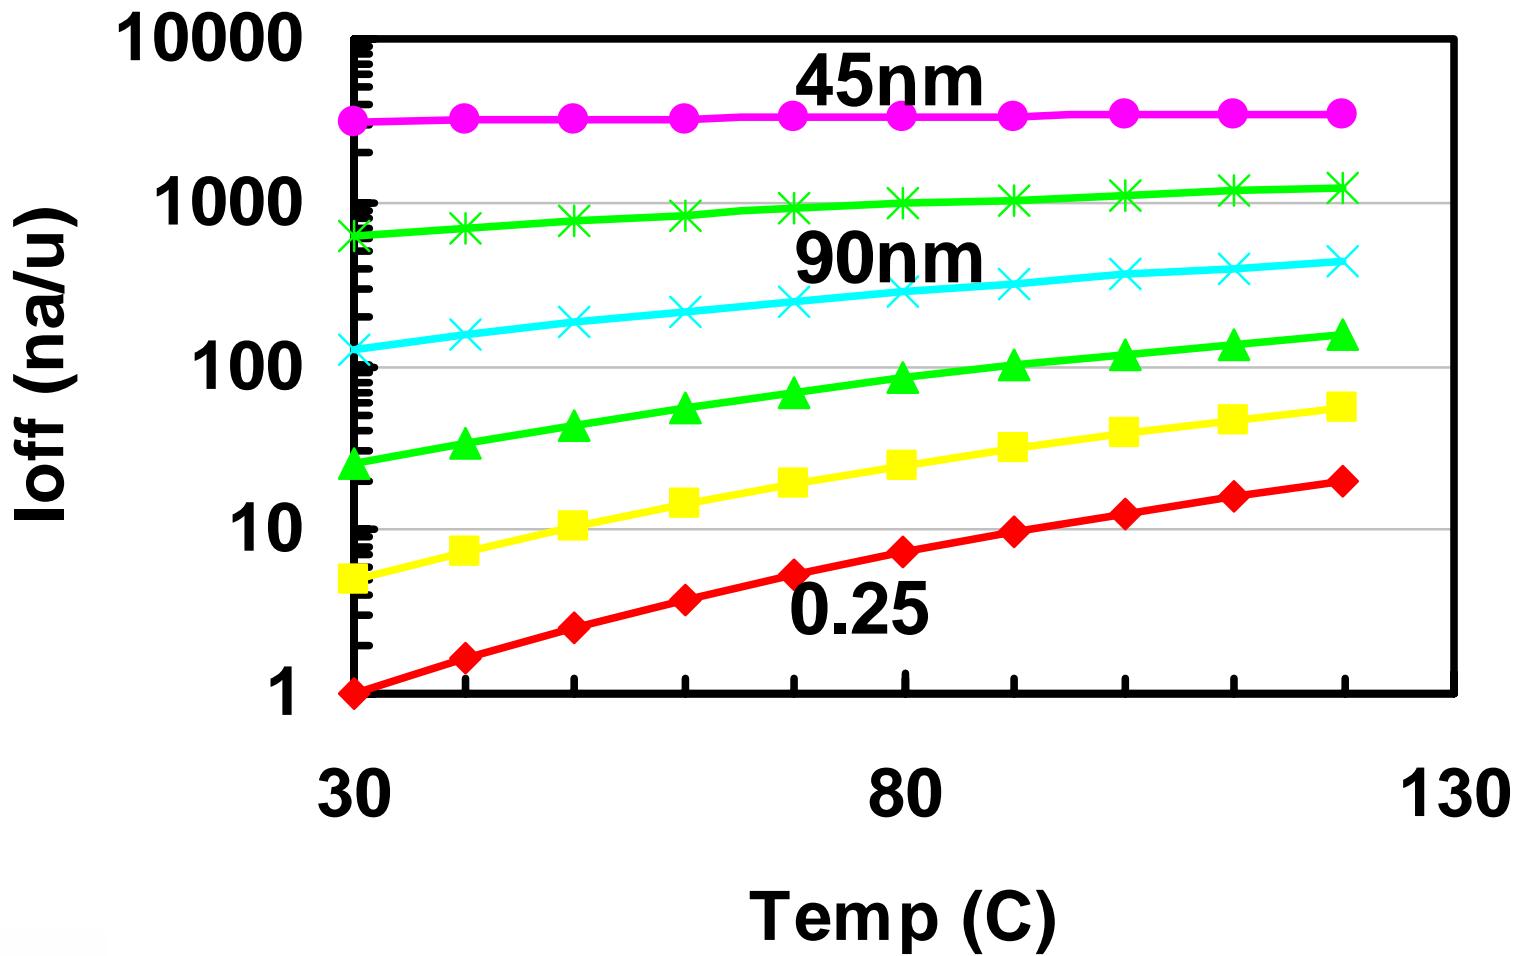
Shekhar Borkar

Intel Corp.

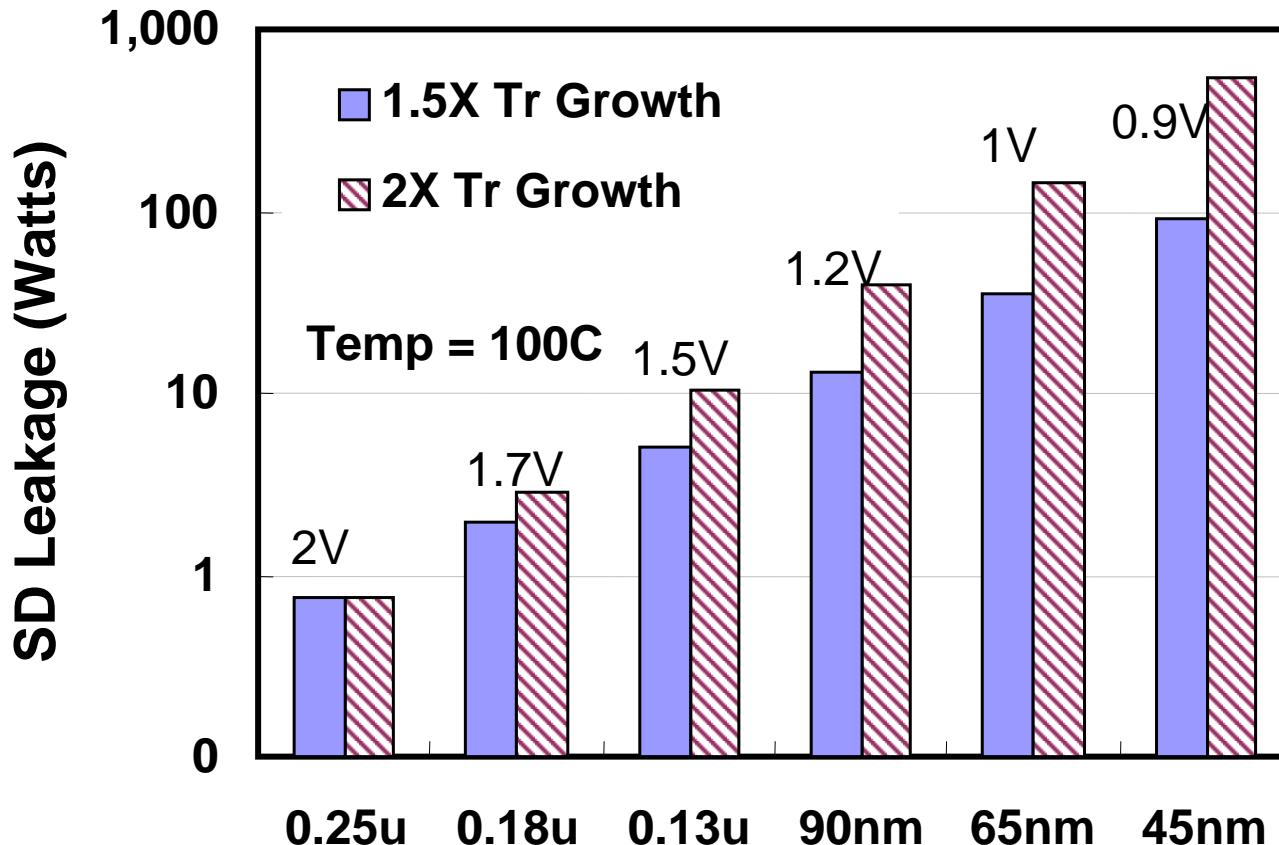
Outline

- **Low Power—Outlook & Challenges**
- **Circuit solutions for leakage avoidance, control, & tolerance**
- **Microarchitecture for Low Power**
- **System considerations**
- **Technology, Circuits, Architecture, and System co-optimization**

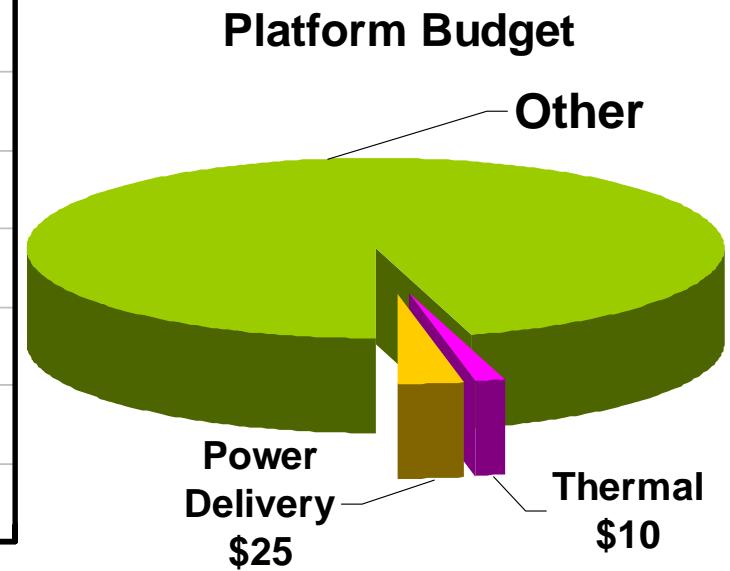
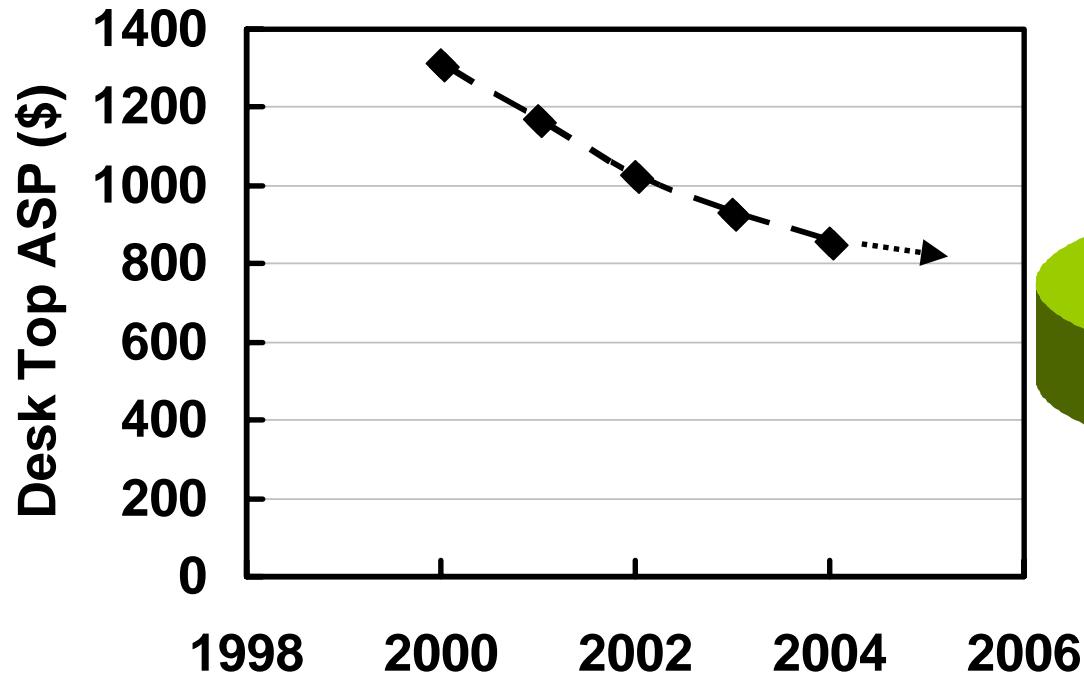
Unconstrained SD Leakage



Unconstrained SD Leakage Power

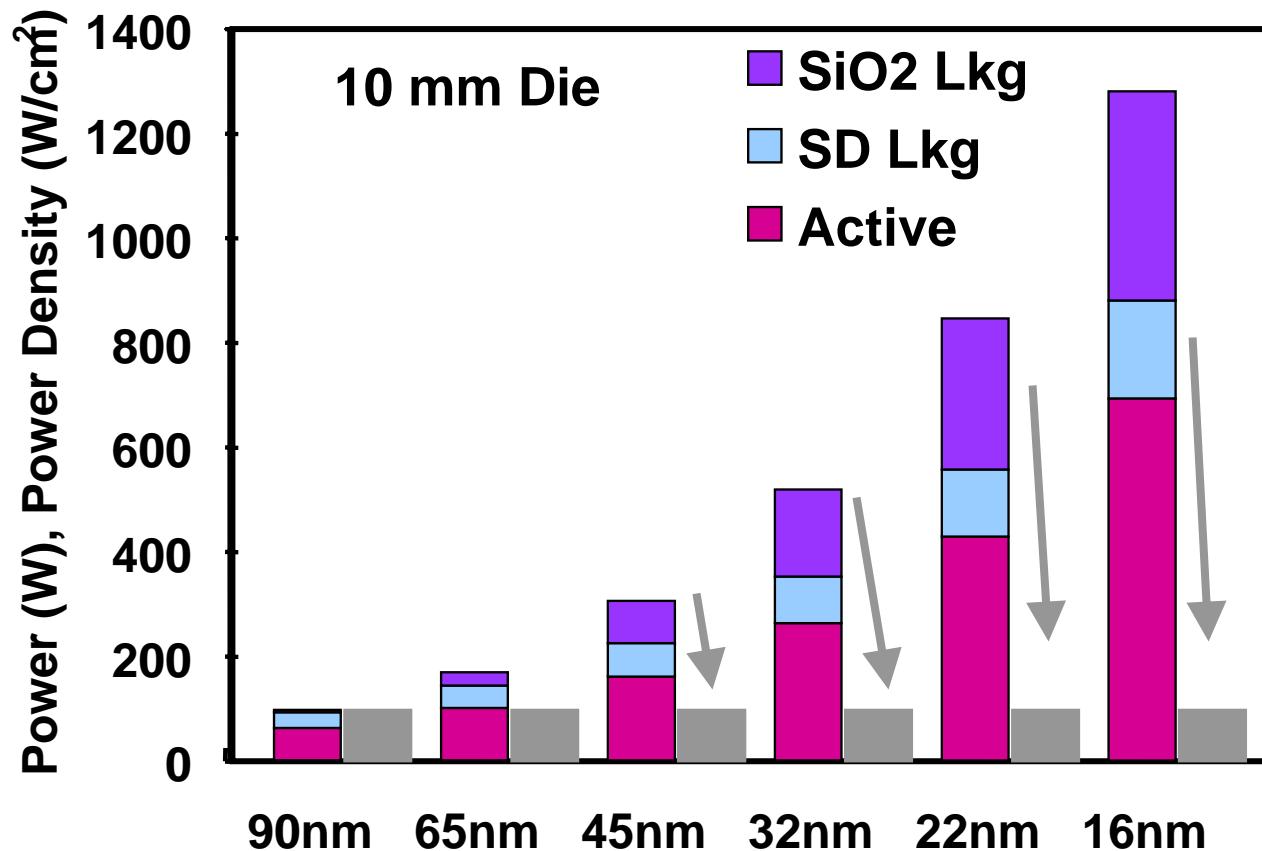


Product Cost Pressure



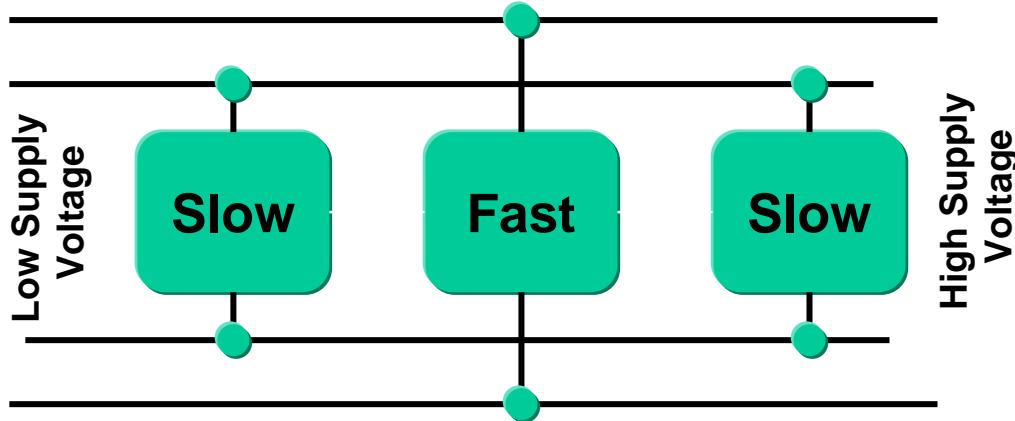
Shrinking ASP & budget for power

Unconstrained Total Power



Technology, Circuits, & Architecture to constrain power

Active Power Reduction



Multiple Supply Voltages

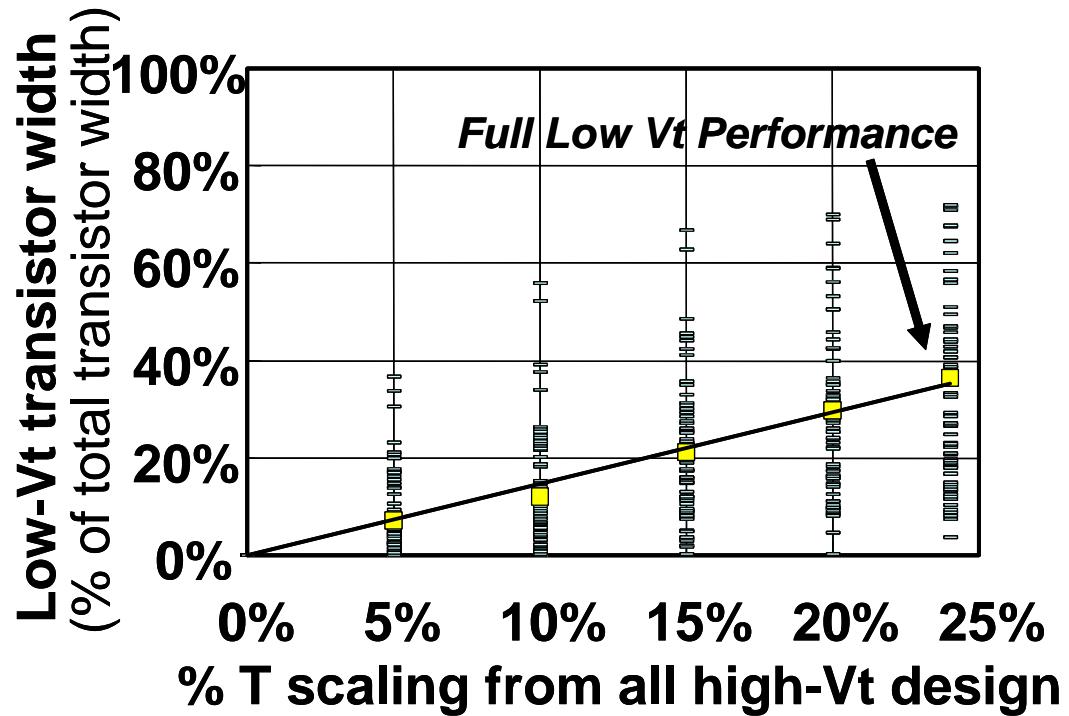
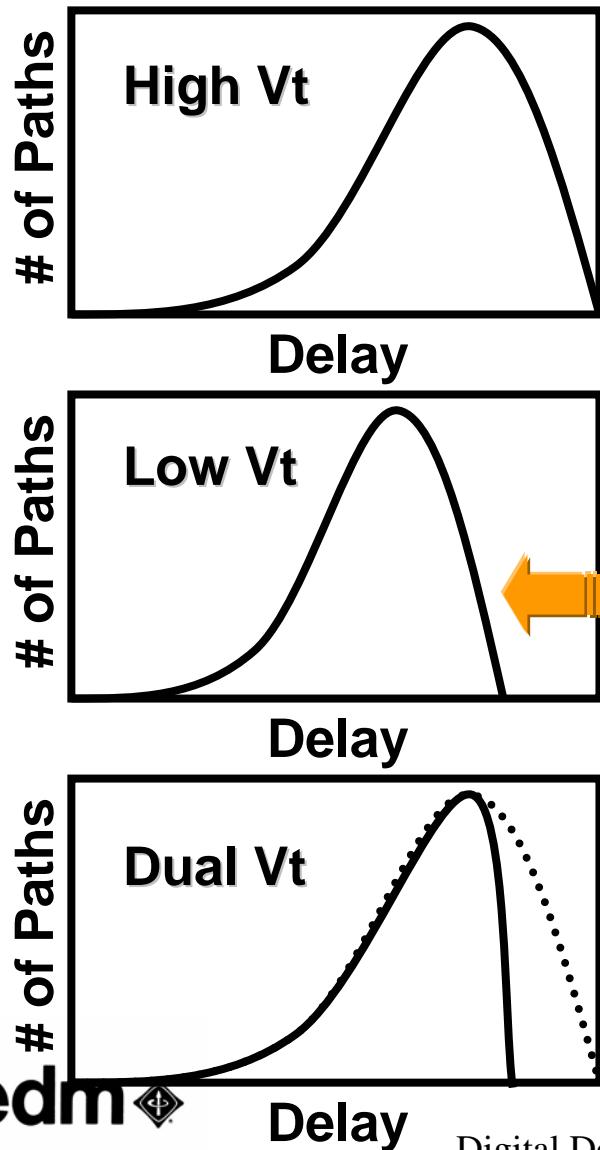
Multiple supply voltages mimic Vdd scaling

Issues:

Performance loss due to interface circuits

Multiple Vdd routing overhead

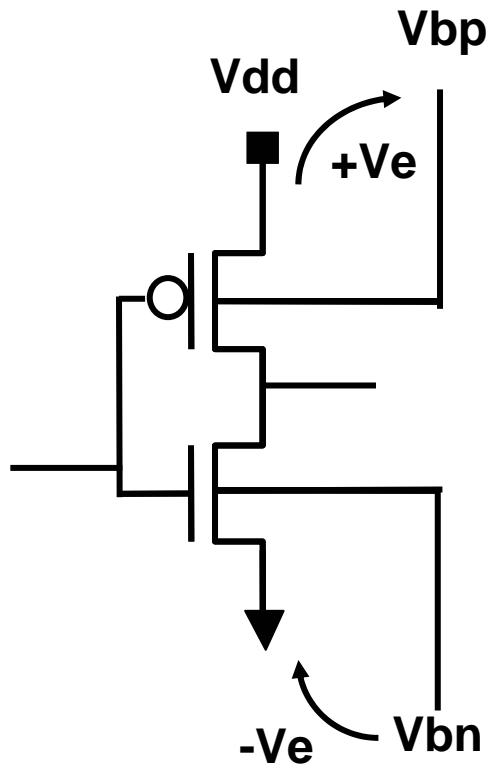
SD Leakage Avoidance—Dual Vt



**Leakage 3X smaller
(Active & Standby)
No performance loss**

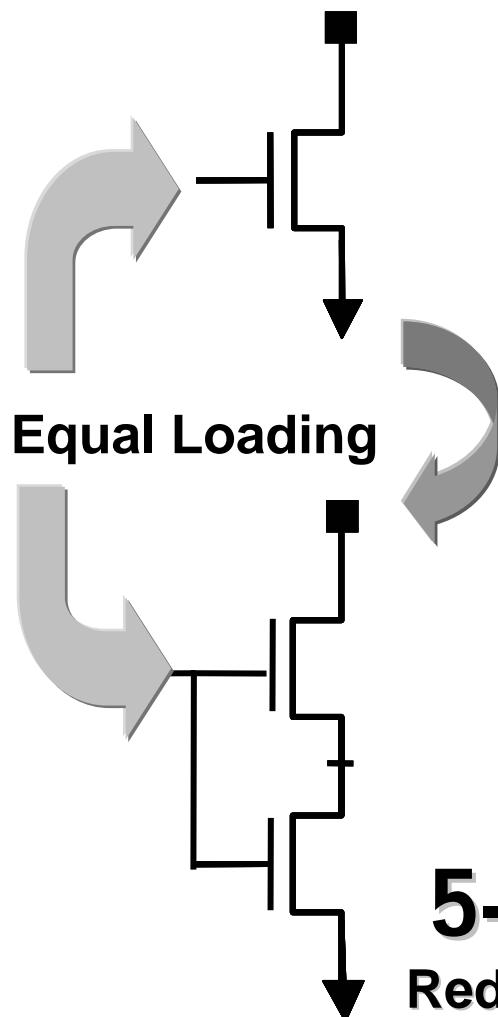
Leakage Control Circuits

Body Bias

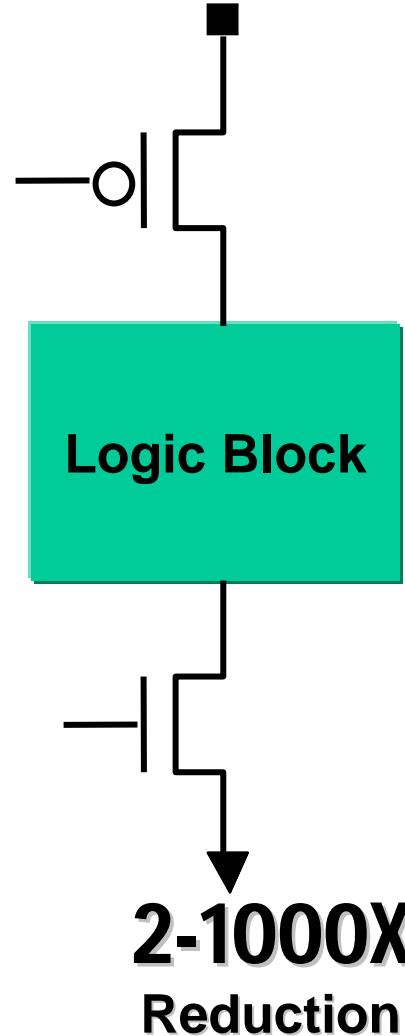


2-10X
Reduction

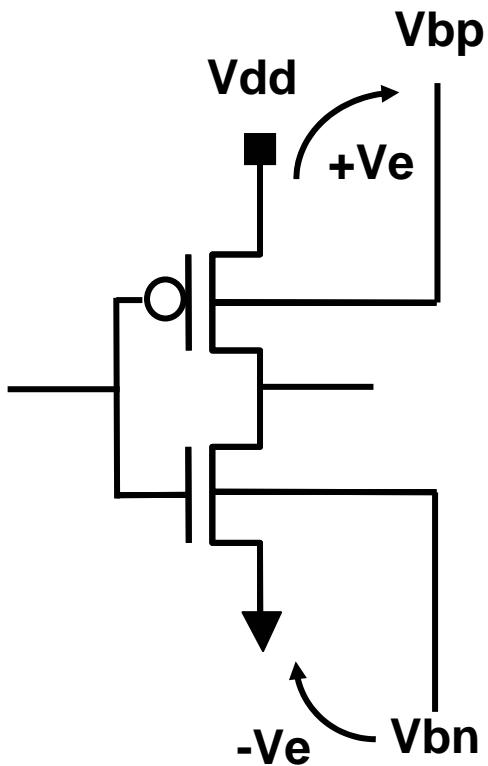
Stack Effect



Sleep Transistor



Body Bias

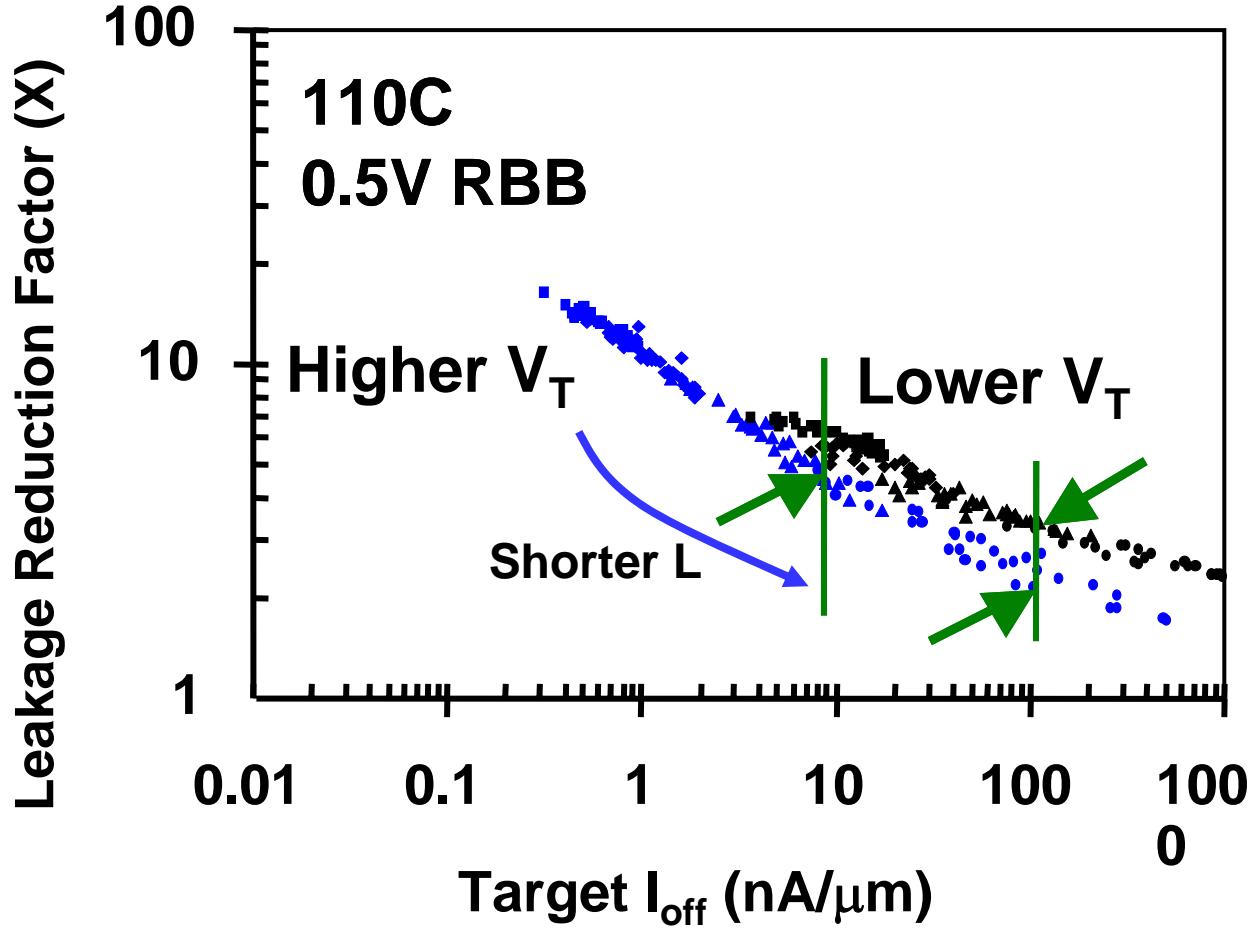


Reverse BB	Forward BB
Increases V_t Reduces I_{off}	Reduces V_t Increases I_{off}
To inactive circuits to reduce leakage	To active circuits to improve performance

Tradeoff & analysis required

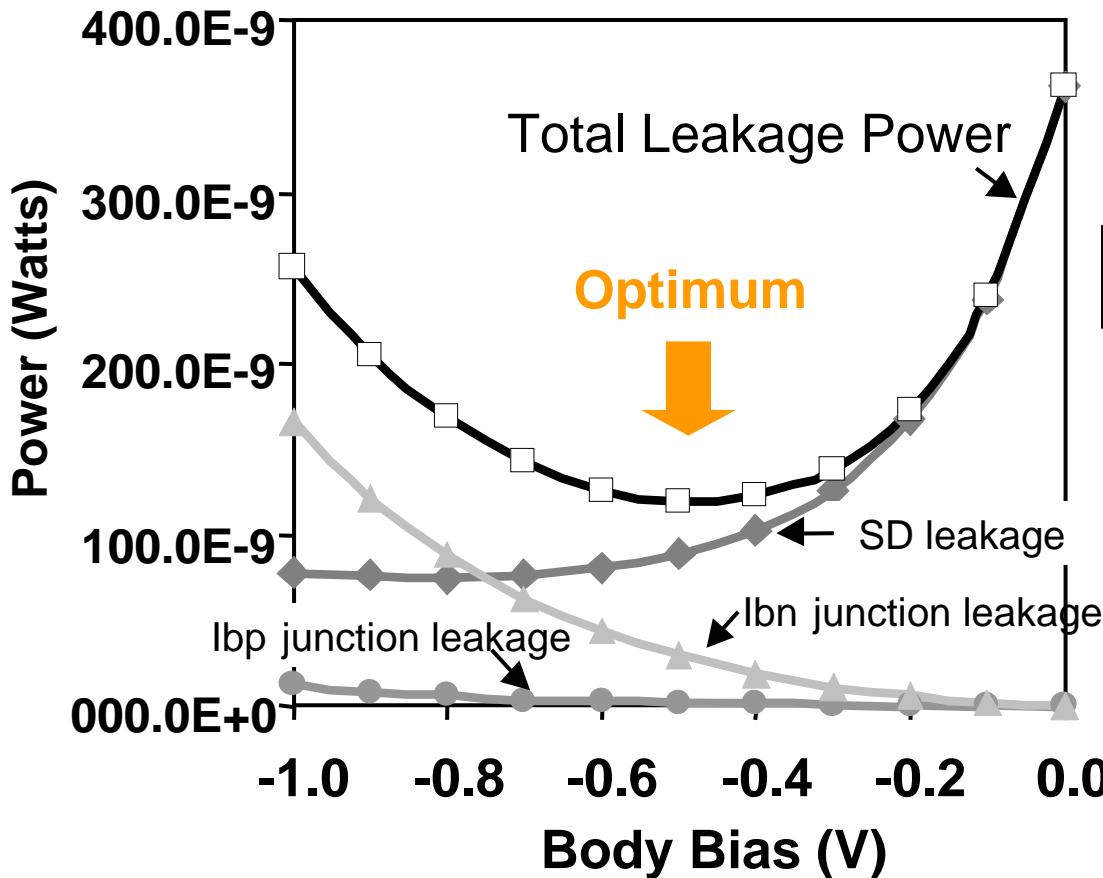
- Applying BB consumes active power
- Switching delay between body bias

Reverse Body Bias



RBB less effective at shorter L and lower V_t

Scalability of RBB



Total Leakage Power
Measured on Test Chip



0.18 μm

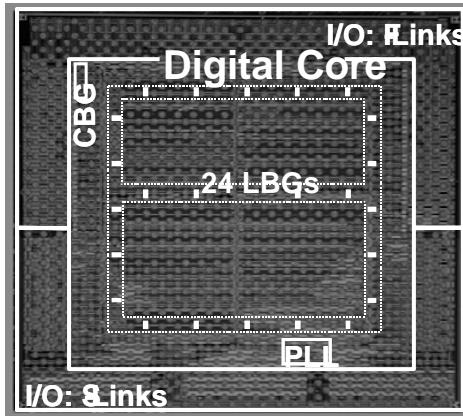
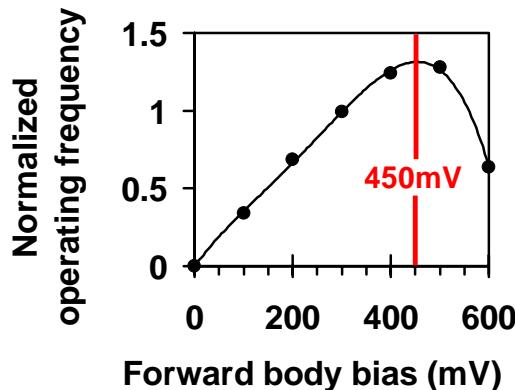
	0.35 μm	0.18 μm
Optimum RBB	2V	0.5V
Ioff Red.	1000X	10X

RBB Less effective with technology scaling

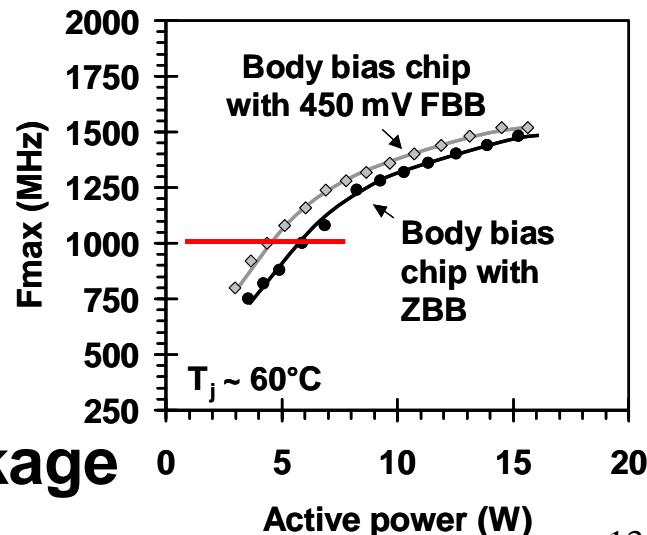
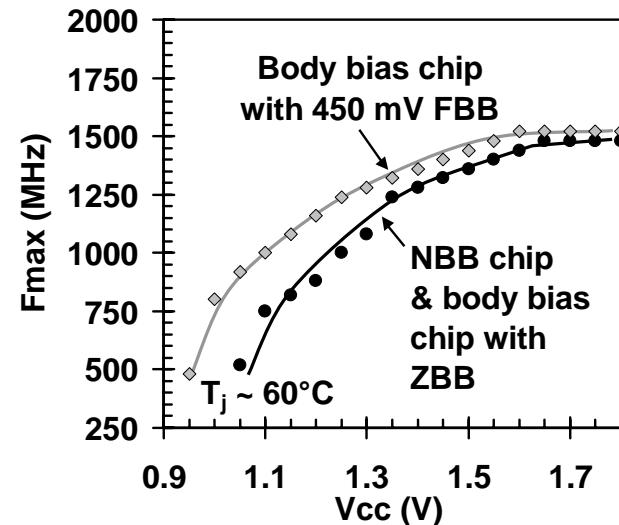
* A. Keshavarzi et. al., 1999 Int. Symp. Low Power Electronics & Design (ISLPED) 12

Forward Body Bias

Router chip with body bias

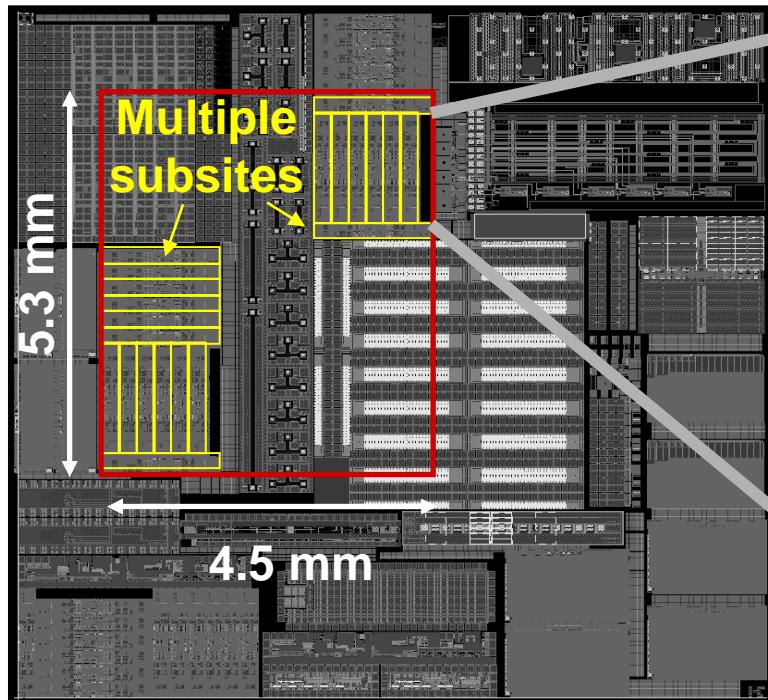


Die size	10.1 X 10.1 mm ²
Technology	150nm CMOS
Transistors	6.6 million
Area overhead	2%
Power overhead	1%



FBB increases frequency & SD leakage

Adaptive Body Bias Experiment

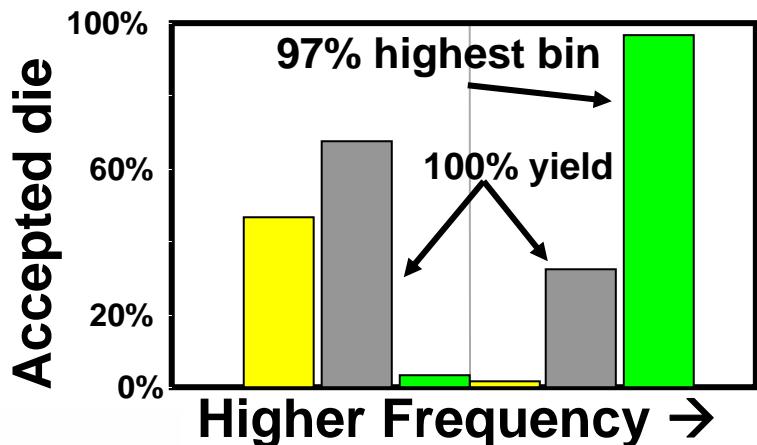
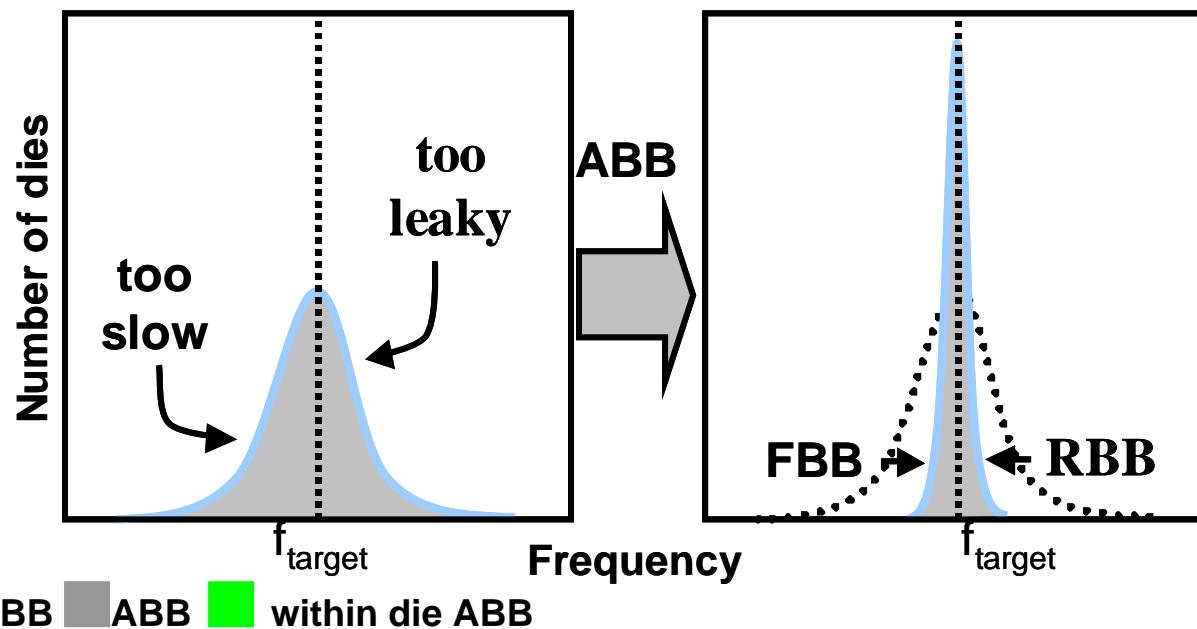


1.6 X 0.24 mm, 21 sites per die
150nm CMOS

Technology	150nm CMOS
Number of subsites per die	21
Body bias range	0.5V FBB to 0.5V RBB
Bias resolution	32 mV

Die frequency: $\text{Min}(F_1..F_{21})$
Die power: $\text{Sum}(P_1..P_{21})$

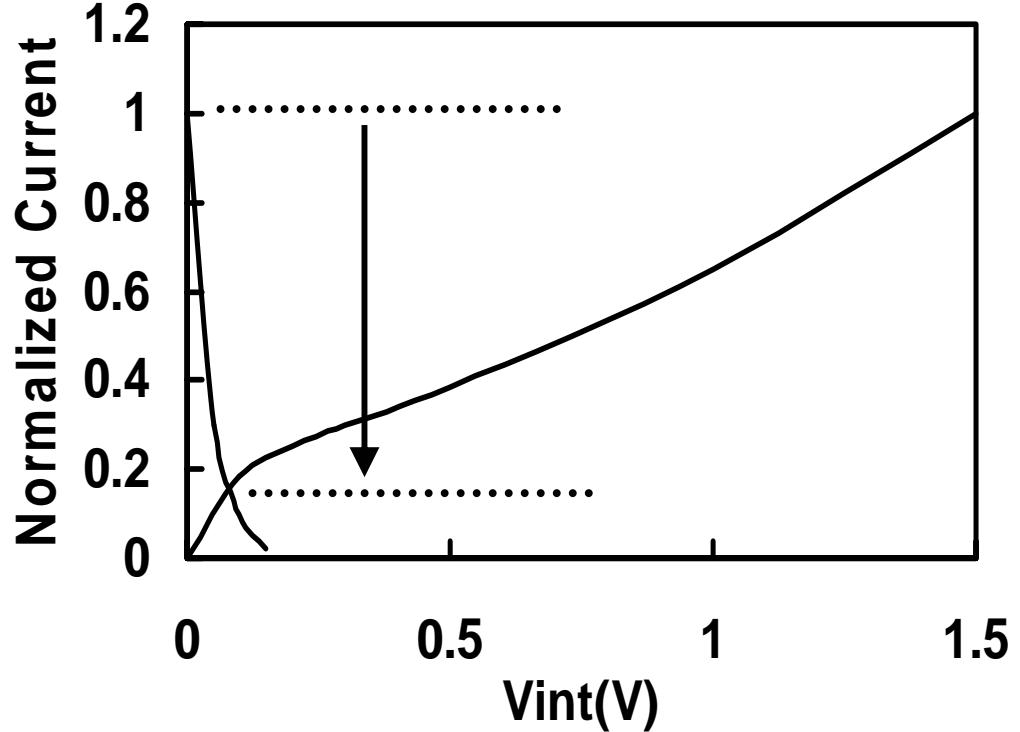
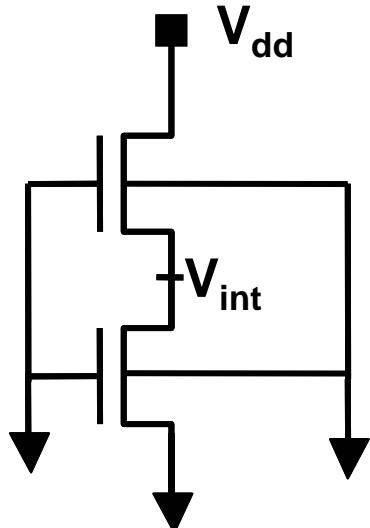
Adaptive Body Bias Results



For given Freq and Power density

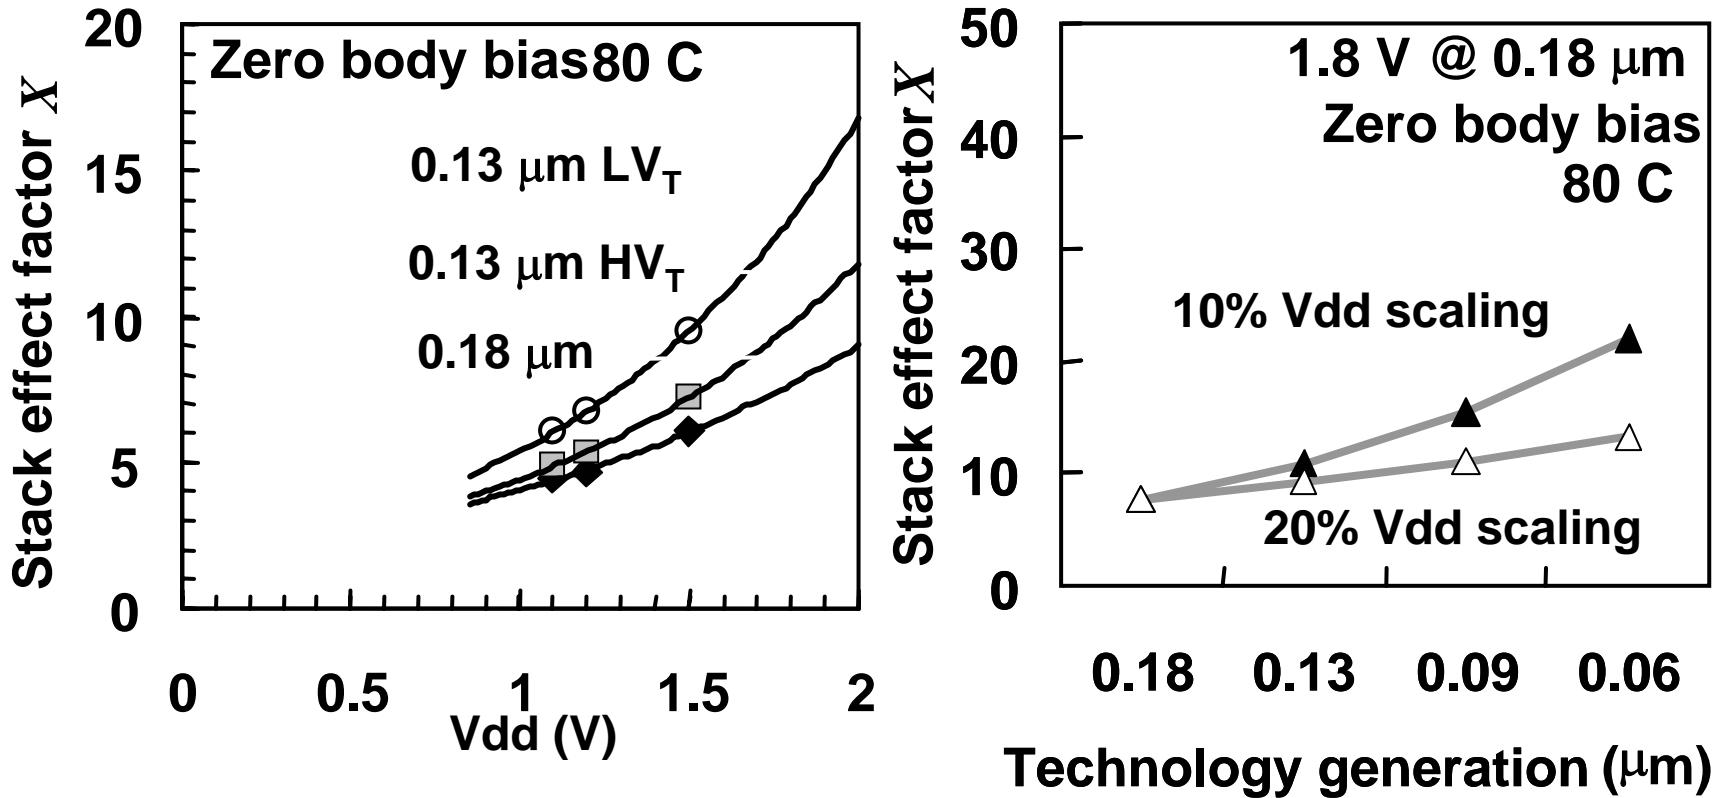
- 100% yield with ABB
- 97% highest freq bin with ABB for within die variability

Transistor Stack



Stack leakage is ~5-10X smaller

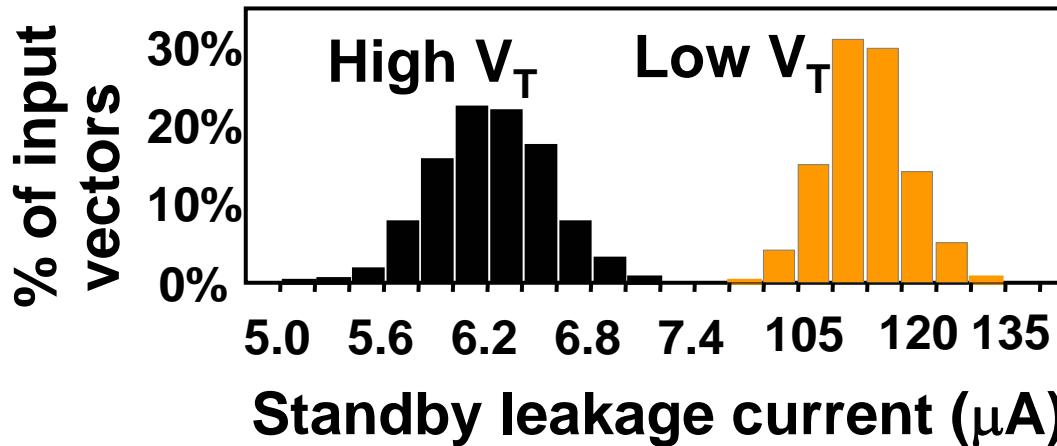
Scalability of Stack Effect



Stack effects becomes stronger with scaling

Exploiting Natural Stacks

32-bit Kogge-Stone adder

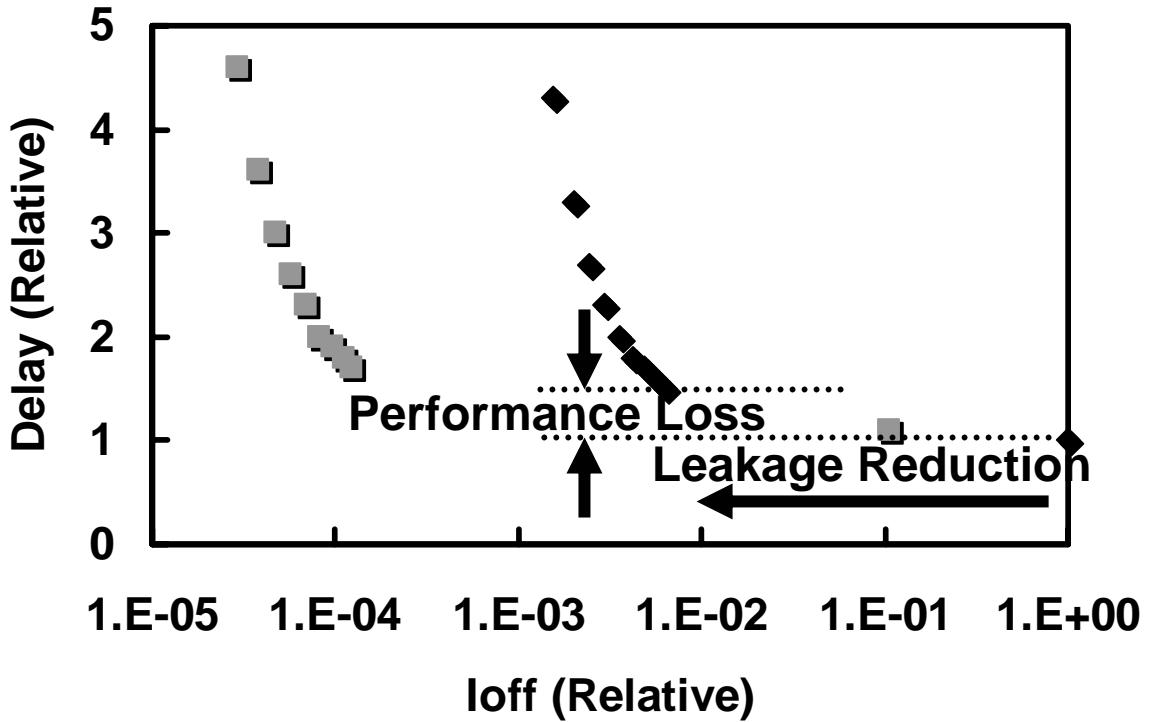
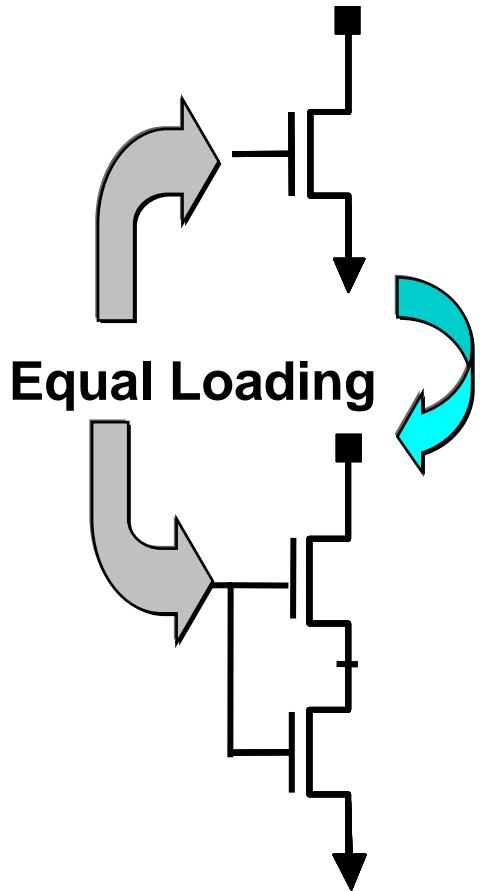


Reduction	Avg	Worst
High V _T	1.5X	2.5X
Low V _T	1.5X	2X

	High V _T	Low V _T
Energy Overhead	1.64 nJ	1.84 nJ
Savings	2.2 μA	38.4 μA
Min time in Standby	84 μS	5.4 μS

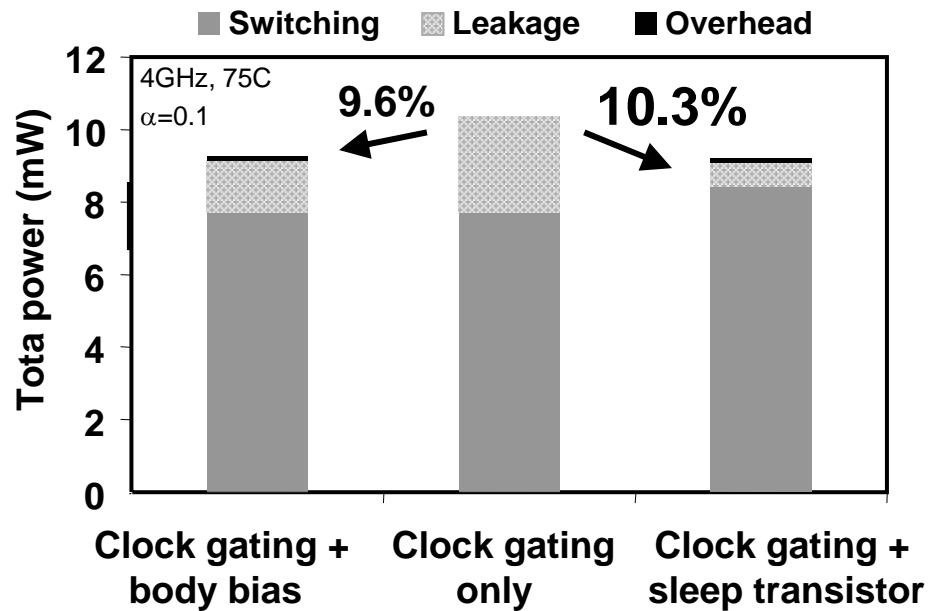
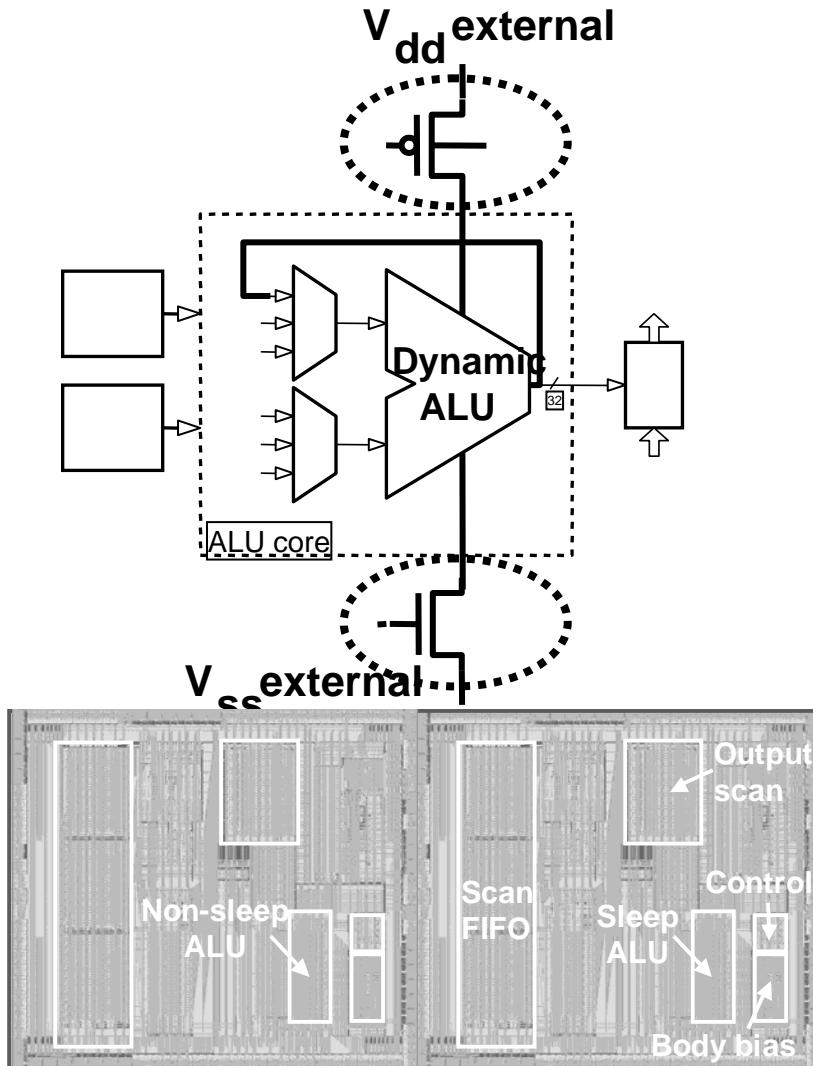
* Y. Ye et. al., 1998 Symp. VLSI Circuits

Stack Forcing



Circuit technique provides additional V_t's

Sleep Transistor Technique

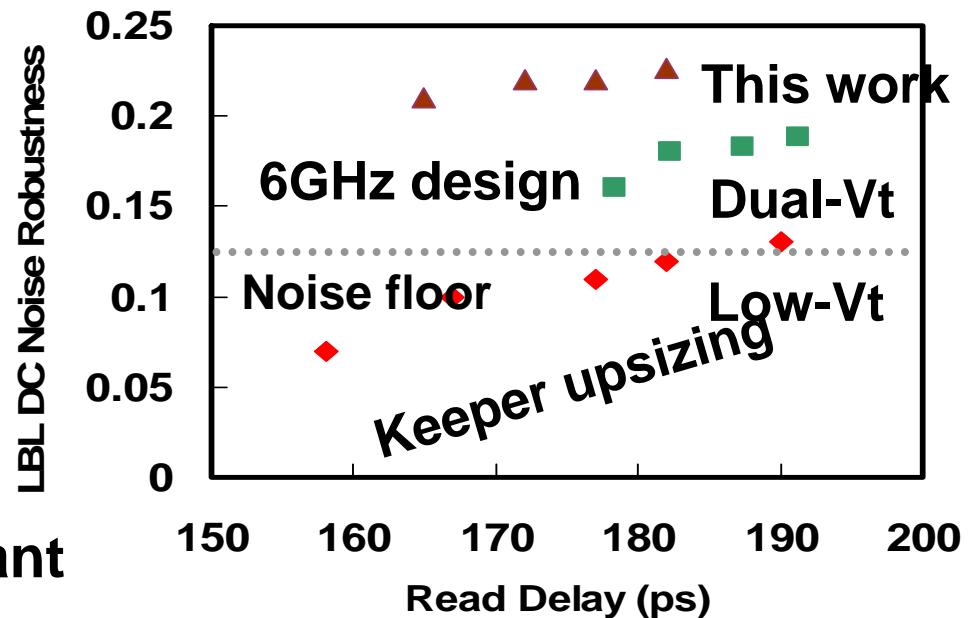
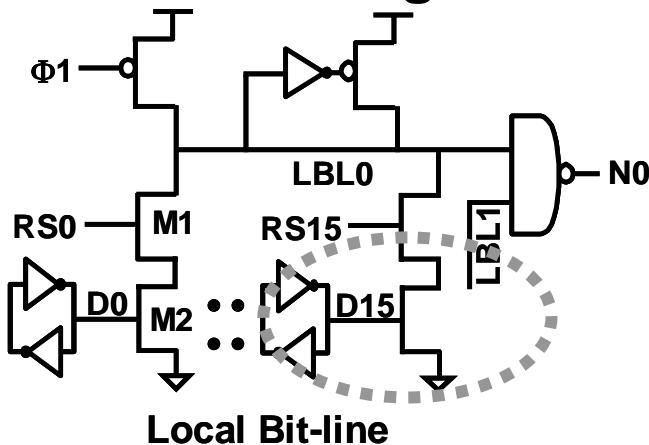


	Frequency change	Leakage savings	Area overhead
PMOS sleep transistor	-2%	13X	11%
PMOS body bias	0%	2X	2%

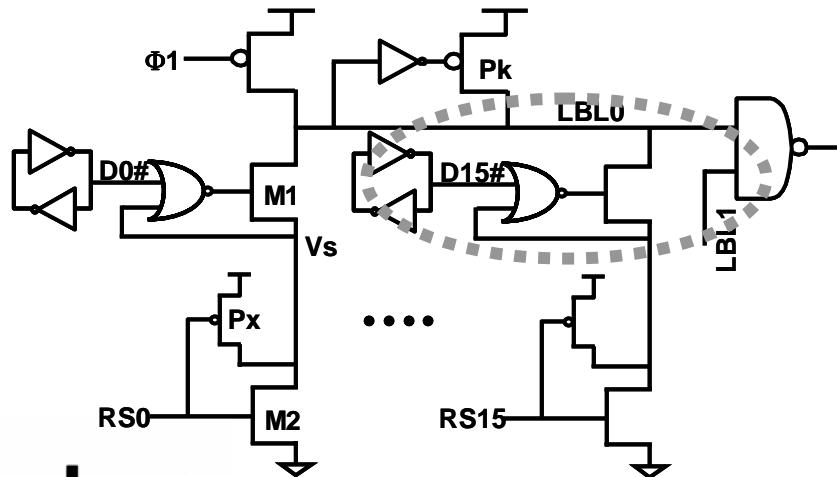
J. Tschanz et al, ISSCC 2003, 6.1

Leakage Tolerant Register File

Domino: leakage-sensitive



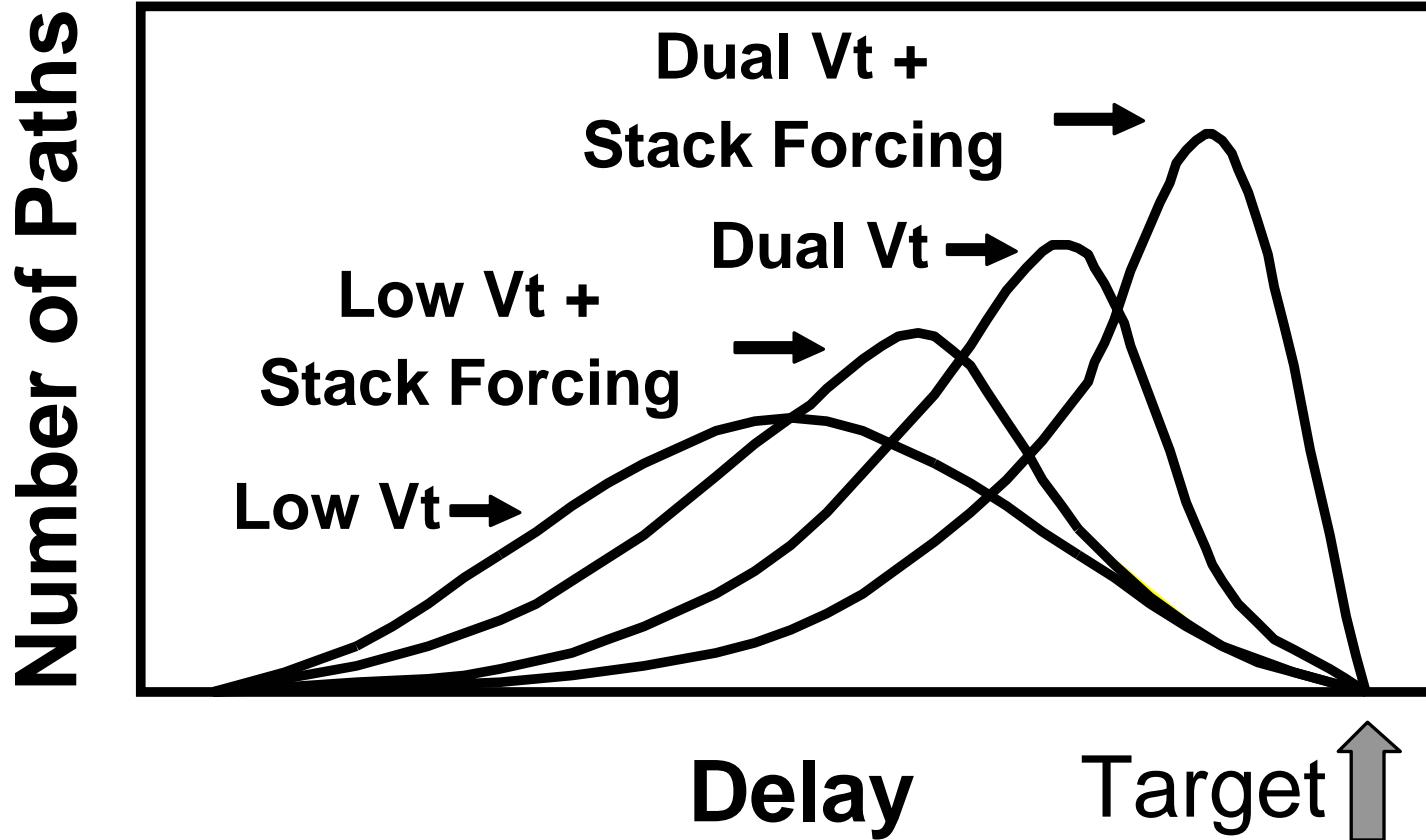
Pseudo-static: leakage-tolerant



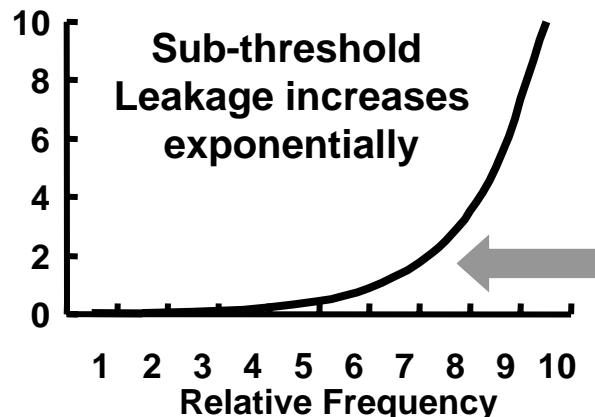
Improves delay
Improves noise margin
Leakage tolerant

* R. Krishnamurthy et. al., 2001 Symp. VLSI Circuits

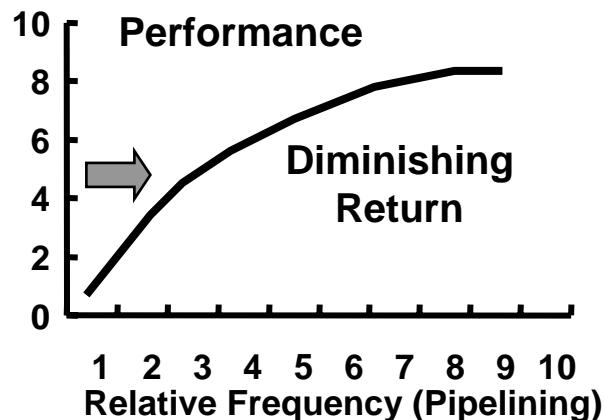
Employing Leakage Control



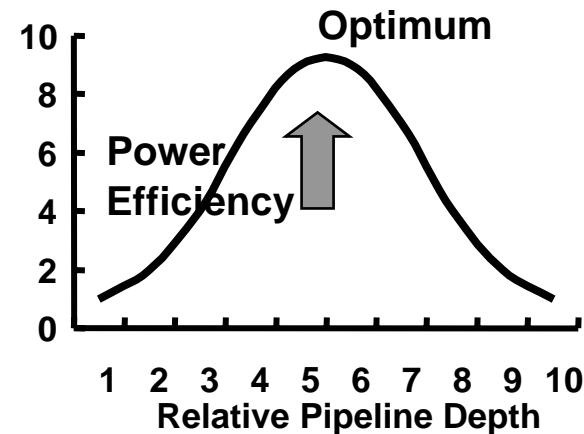
Optimum Frequency



Process Technology



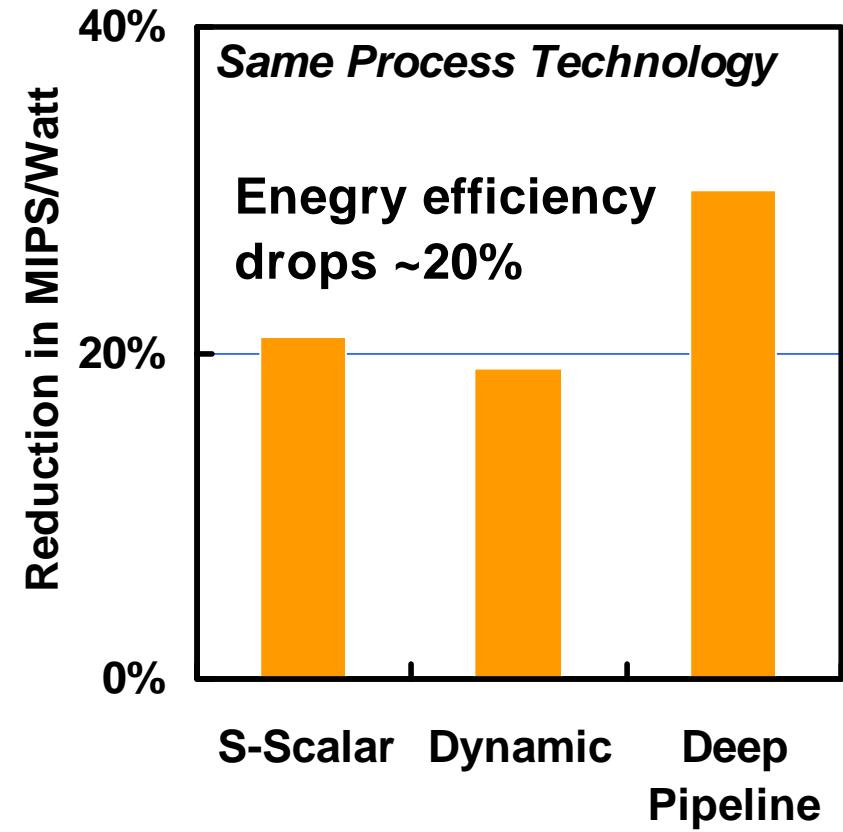
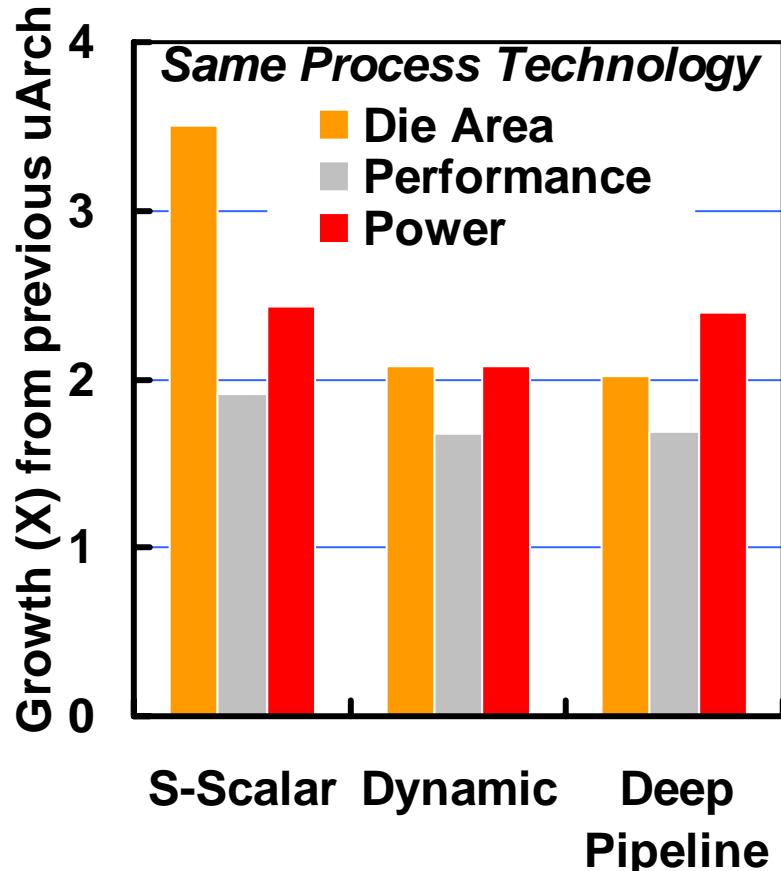
Pipeline & Performance



Pipeline Depth

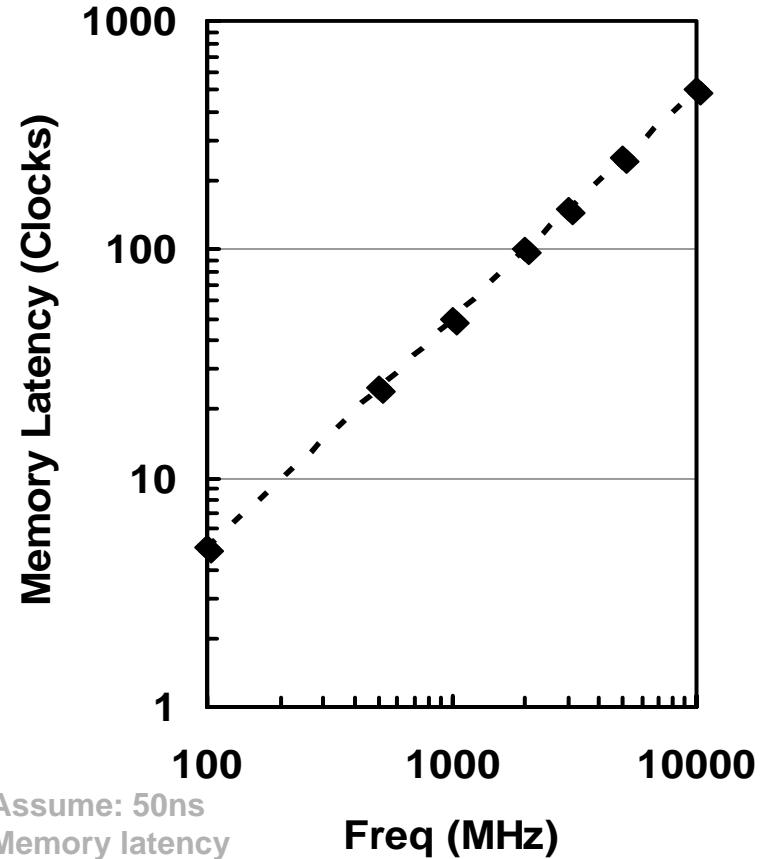
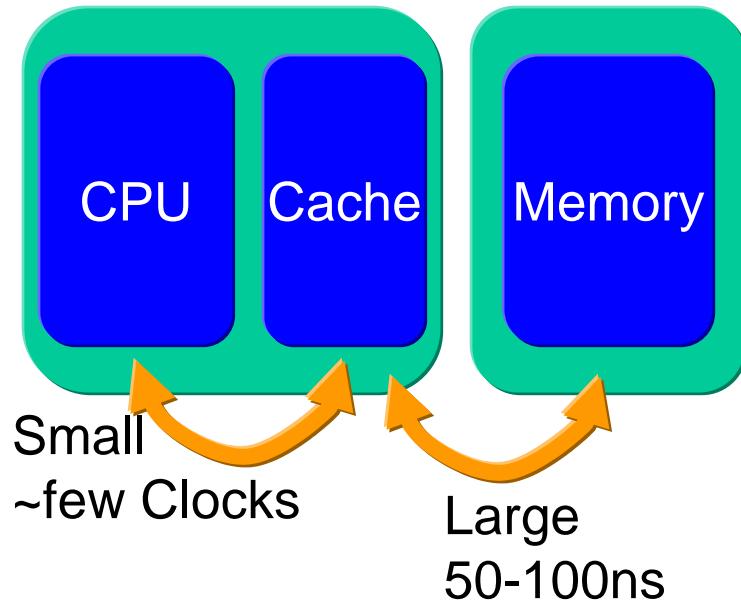
Maximum performance with
• Optimum pipeline depth
• Optimum frequency

Efficiency of Microarchitectures



Employ efficient microarchitecture and design

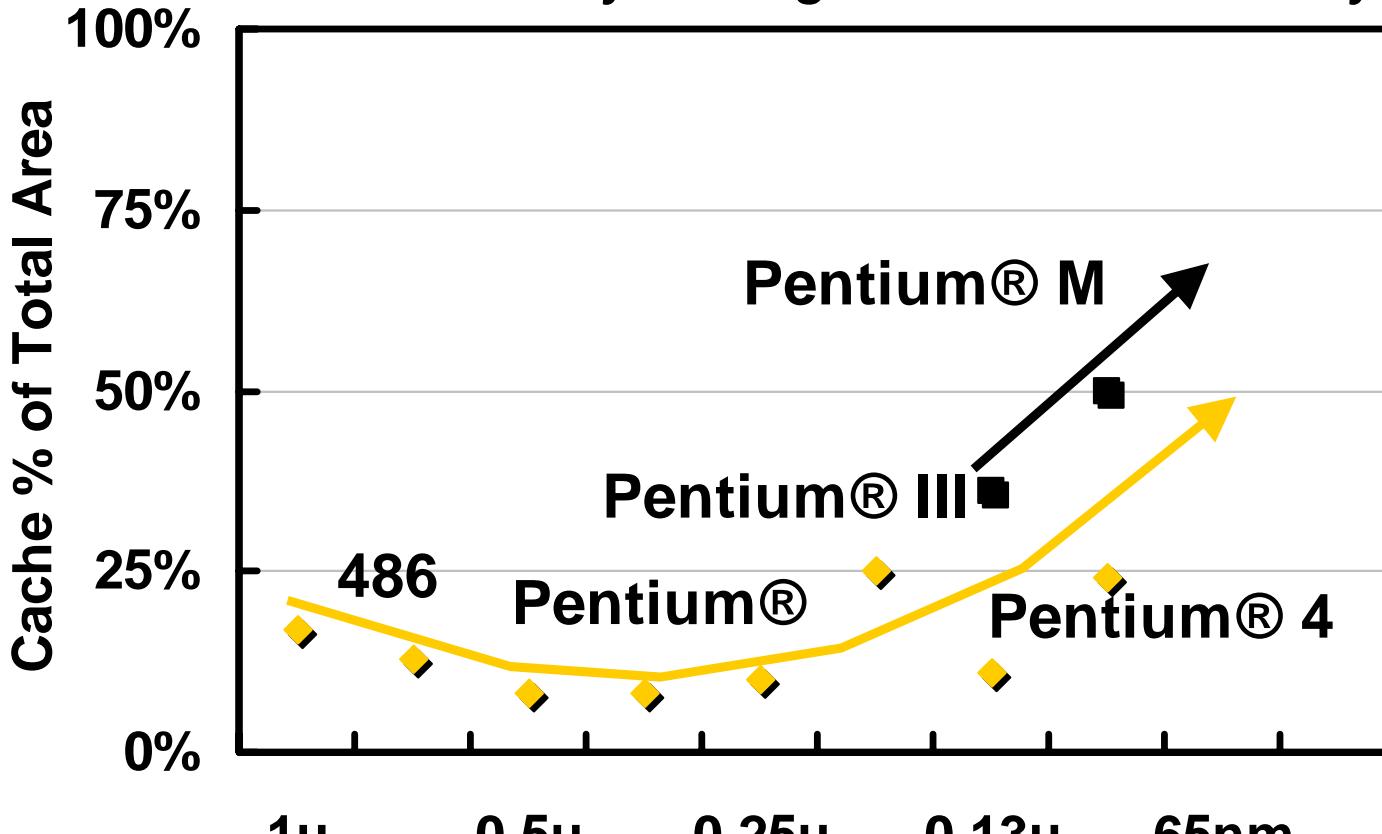
Memory Latency



**Cache miss hurts performance
Worse at higher frequency**

Increase on-die Memory

Power Density of Logic $\approx 10X$ of Memory

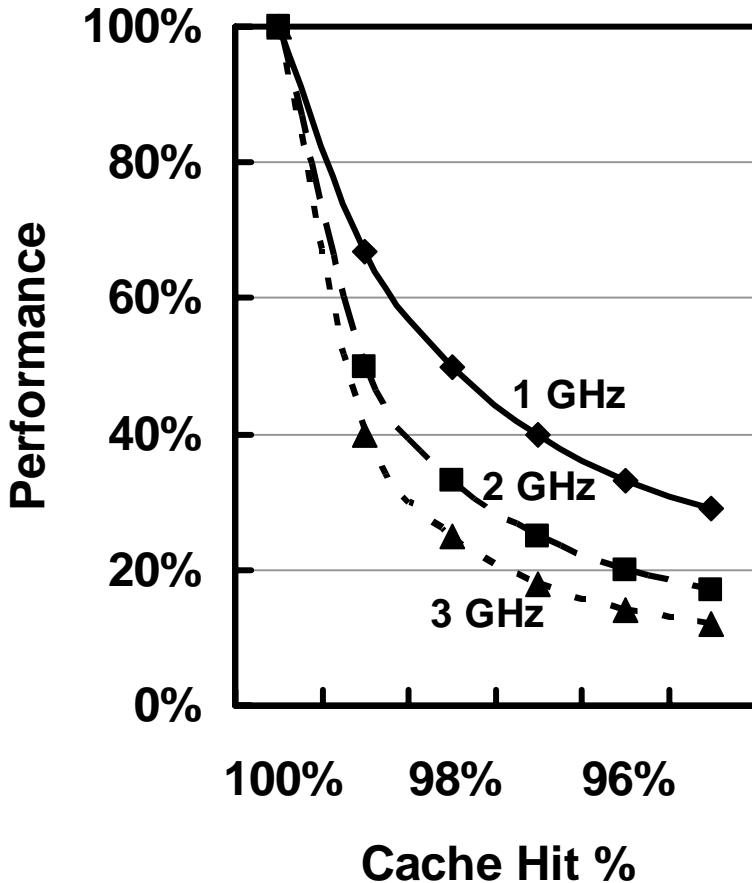


1u 0.5u 0.25u 0.13u 65nm

Increased Data Bandwidth & Reduced Latency

Hence, higher performance for much lower power

Multi-threading



*Thermals & Power Delivery
designed for full HW utilization*

Single Thread

Full HW Utilization

ST **Wait for Mem**

Multi-Threading

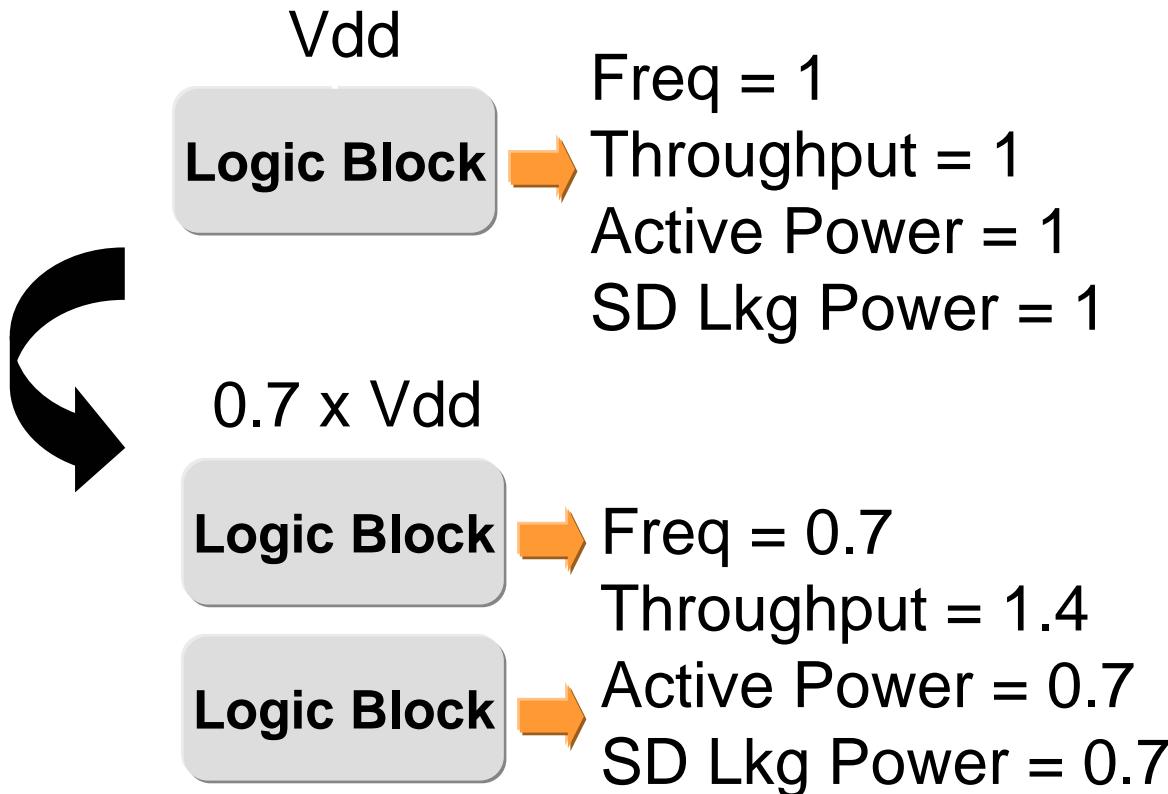
MT1 **Wait for Mem**

MT2 **Wait**

MT3

**Multi-threading improves performance
without impacting thermals & power delivery**

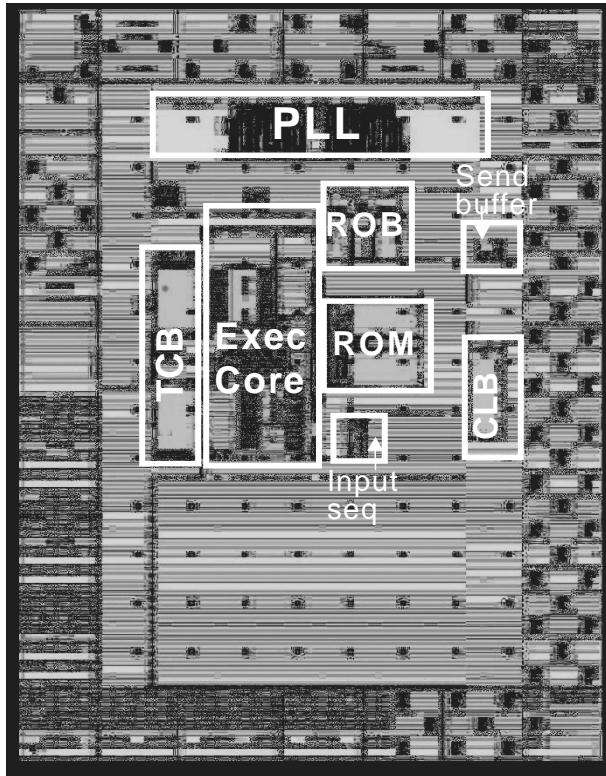
Throughput Oriented Design



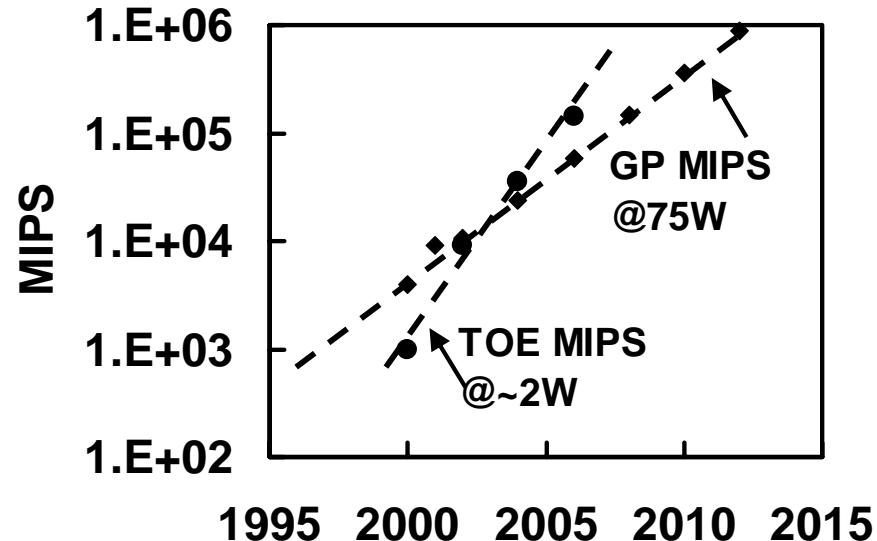
Higher logic throughput, yet lower power

Special Purpose Hardware

TCP Offload Engine



2.23 mm X 3.54 mm, 260K transistors



Opportunities:
Network processing engines
MPEG Encode/Decode engines
Speech engines

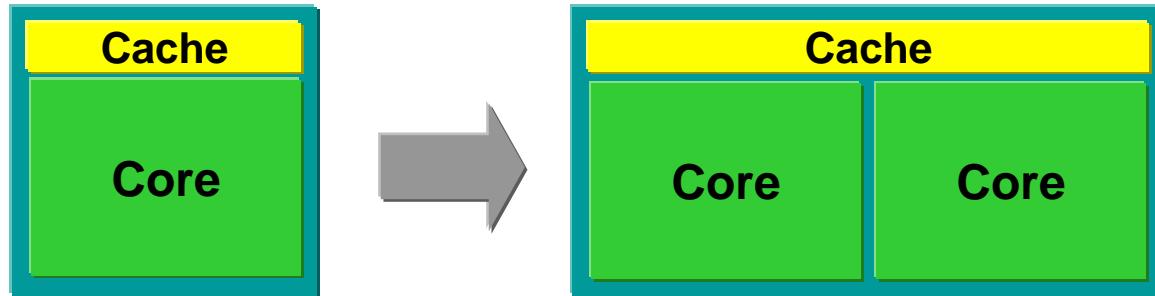
Special Purpose HW—Best Mips/Watt

Chip Level Multi-processing

Rule of thumb:

Voltage	Frequency	Power	Performance
1%	1%	3%	0.66%

In the same process technology...



$$\text{Voltage} = 1$$

$$\text{Freq} = 1$$

$$\text{Area} = 1$$

$$\text{Power} = 1$$

$$\text{Perf} = 1$$

$$\text{Voltage} = -15\%$$

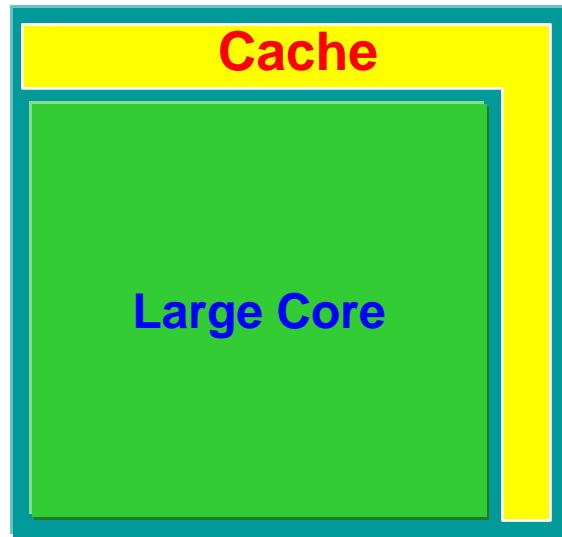
$$\text{Freq} = -15\%$$

$$\text{Area} = 2$$

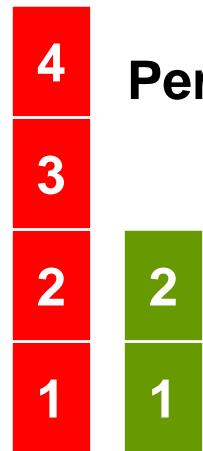
$$\text{Power} = 1$$

$$\text{Perf} = \sim 1.8$$

Multi-Core



Power

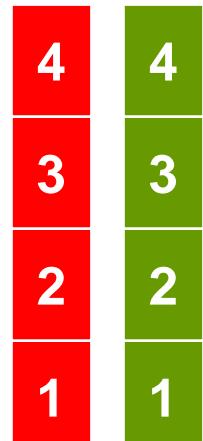
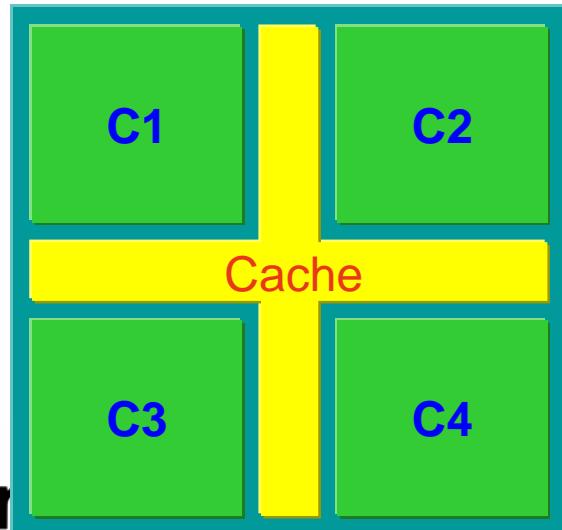


Performance



Power = 1/4

Performance = 1/2

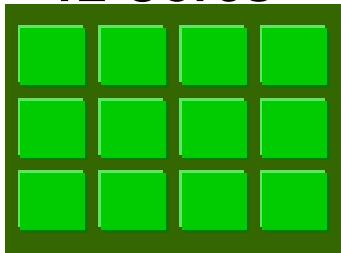


Multi-Core:
Power efficient
Better power and
thermal management

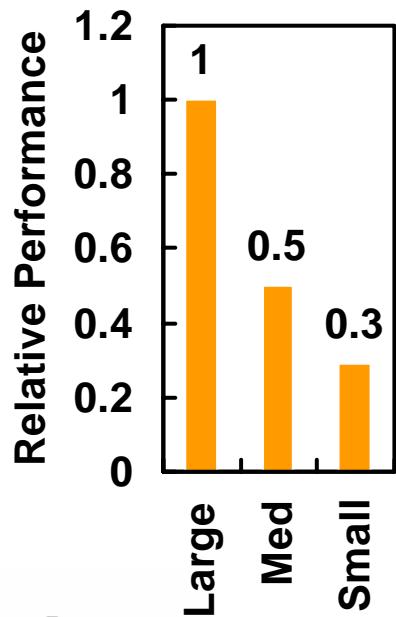
From Multi to Many...

13mm, 100W, 48MB Cache, 4B Transistors, in 22nm

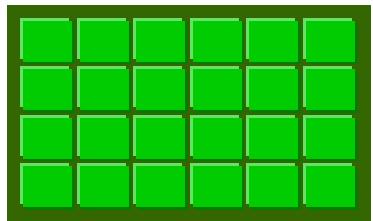
12 Cores



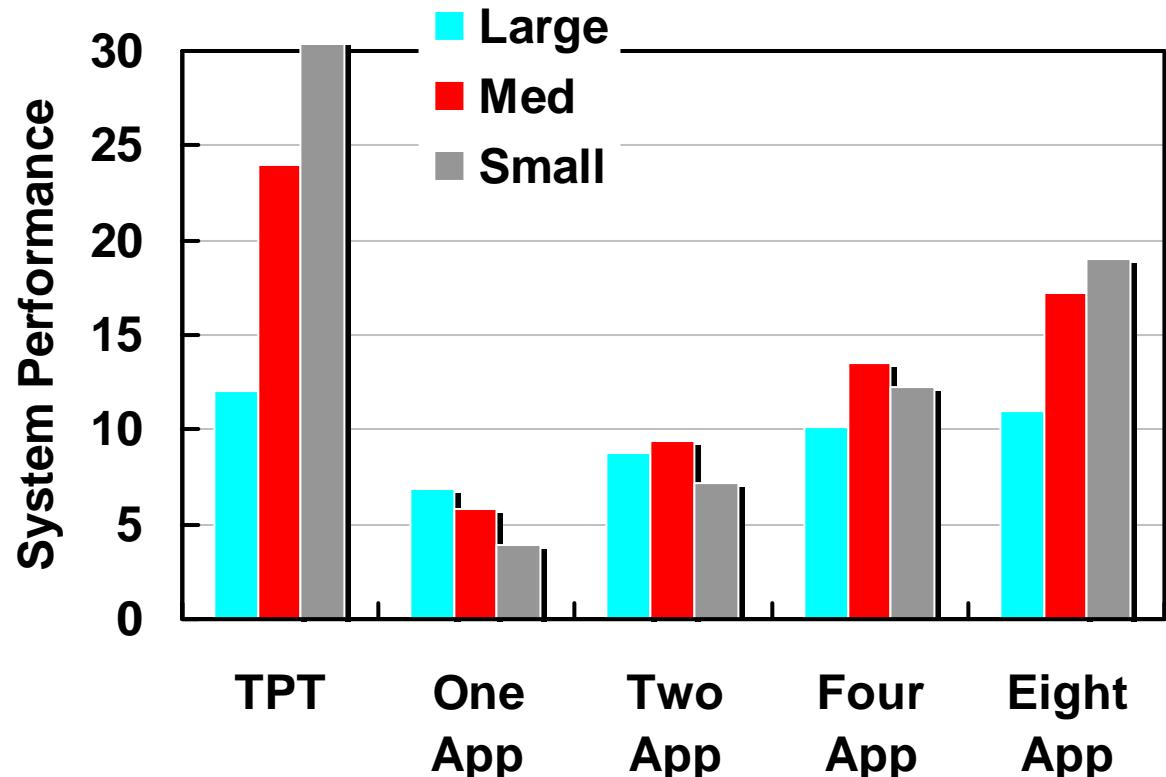
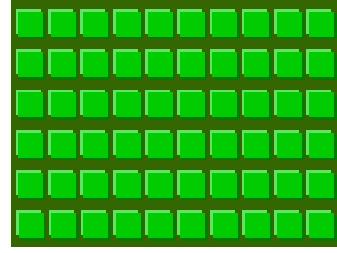
Single Core Performance



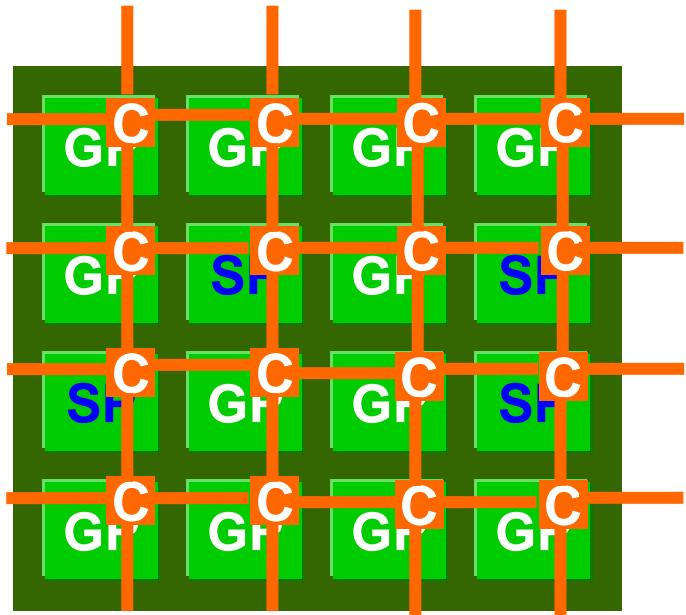
48 Cores



144 Cores



Future Multi-core Platform



General Purpose Cores

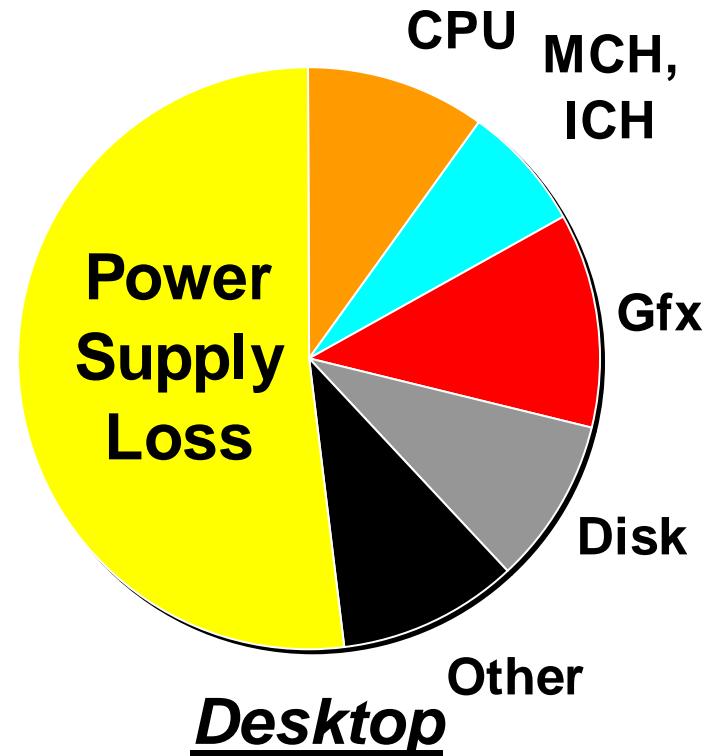
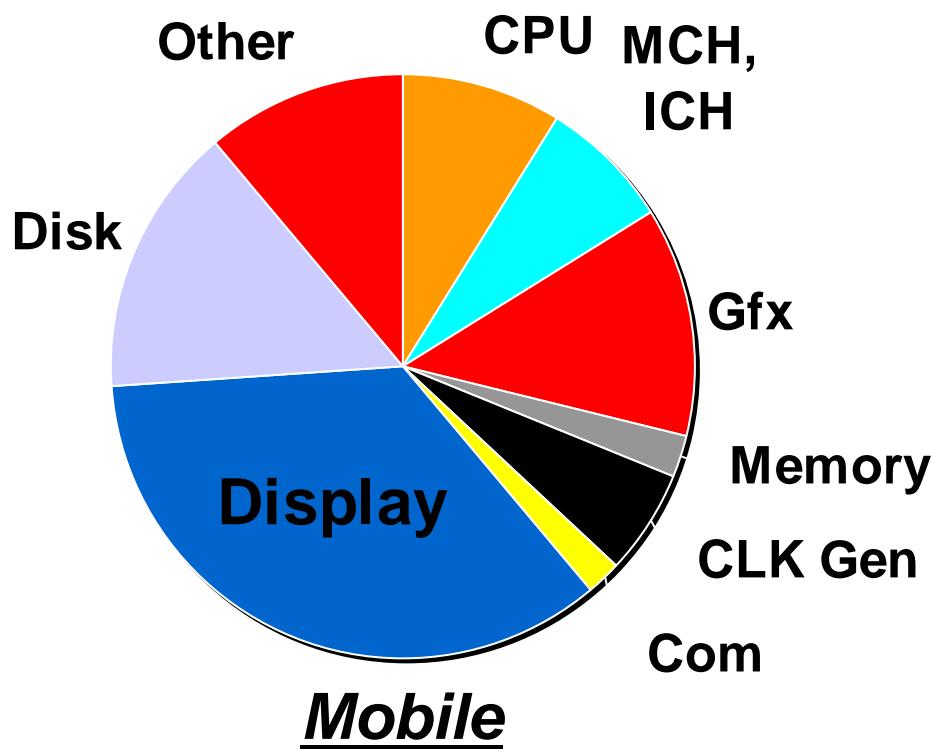
Special Purpose HW

Interconnect fabric

Heterogeneous Multi-Core Platform

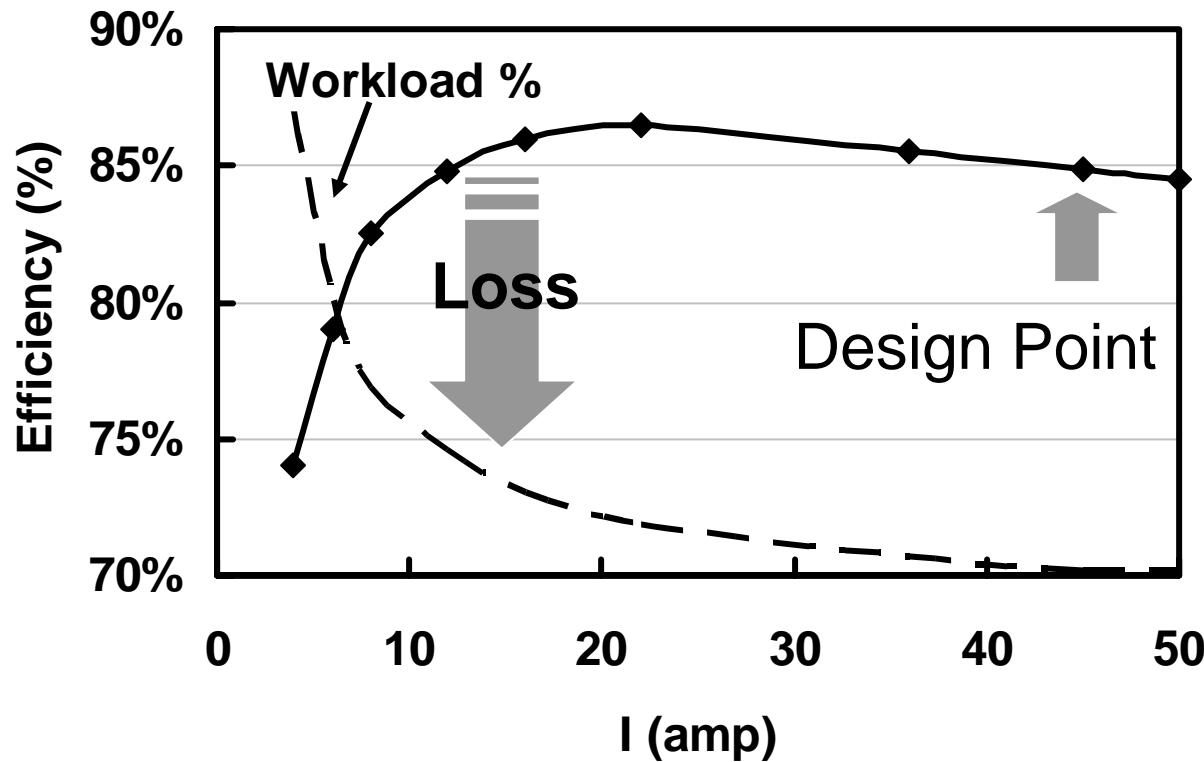
Low Power & Platform

Breakdown of Platform Power



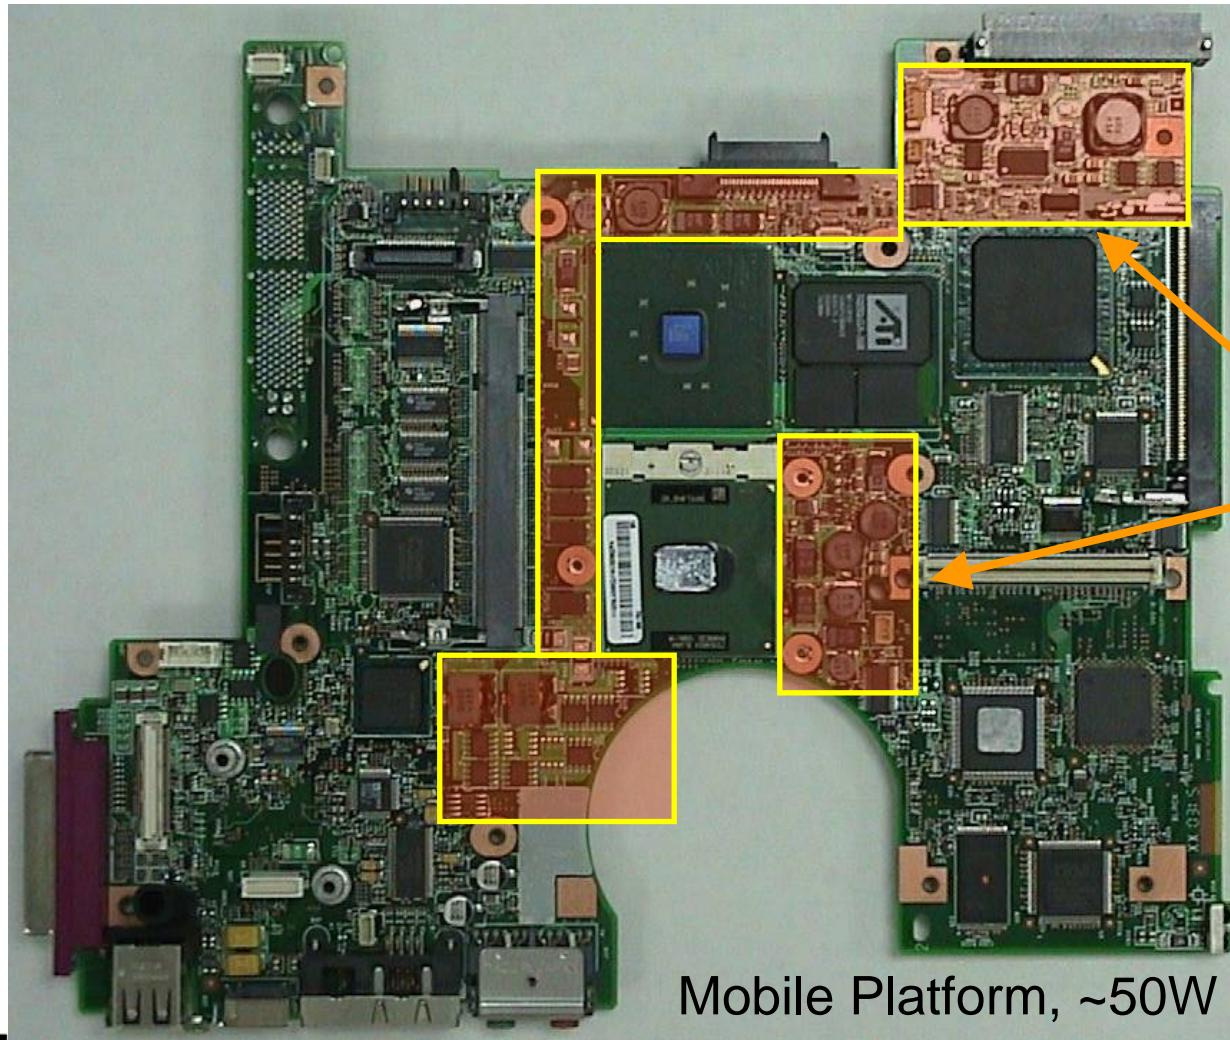
**CPU & electronics power in a platform is low
Must reduce power of other platform ingredients**

Efficiency of Power Delivery



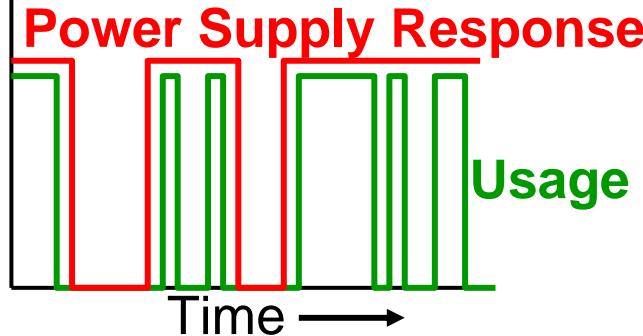
Difficult to maintain high efficiency across workloads
Power supply design needs to exploit new technologies

Typical Power Delivery System

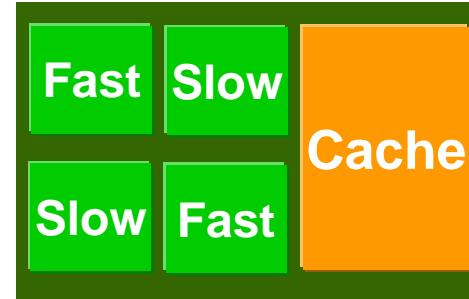


Power
Delivery
Electronics

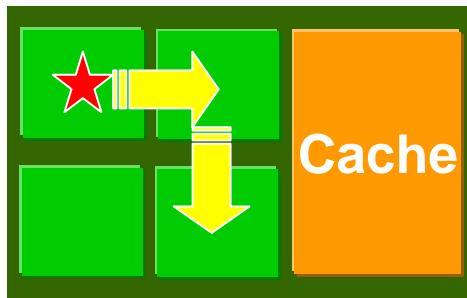
Fine-grain Power Management



Improve response time
Bring power supply closer



Provide multiple supply voltages
Fine grain Vdd and Freq scaling



Core Hopping for hot spot & power density management

Software & Low Power

Applications

Provide guidance to compilers and operating system through policy

Compilers

Add hooks for fine grain power management, and assist OS by providing checkpoints

Operating System

Fine grain power management algorithm
Control HW through firmware

Firmware

Hardware control
Detect conditions, take action, notify OS

Fine grain power management is possible only by cooperation of the entire software stack

Co-optimization

Technology	Circuits	μ Arch	Platform	Software
SD Leakage (loff)	Leakage avoidance, control & tolerance	Optimal pipeline Throughput oriented micro-architecture	Fine grain power management hardware & software	
Number of Vt's				
Gate dielectric (Tox)				
Supply voltage (Vdd)	Multiple Vdd circuits	Partitioning of logic Error correction	Efficient power delivery & management of multiple Vdd's	
Transistor density	Tradeoff frequency for transistors	Throughput oriented micro-architecture	Software and applications to utilize logic throughput	
Interconnect	RC delay tolerance		Interconnect cost vs efficiency of power delivery	

Summary

- **Power and energy will limit integration**
- **Employ circuit techniques for SD leakage avoidance, control, & tolerance**
- **Microarchitecture for Low Power**
- **Opportunities to save power at the platform level—Hardware & Software**
- **Technology, Circuits, Architecture, and System co-optimization will be key**