

Scaling the Power Wall— Revisiting the Low-Power Design Rules



IEEE Santa Clara Valley Chapter SSCS
November 15, 2007

Jan M. Rabaey

Director Gigascale Systems
Research Center (GSRC)

Co-Director Berkeley Wireless
Research Center (BWRC)

University of California, Berkeley

Low Power Design Rules (Anno 1996)

The Low-Energy Roadmap

- Voltage as a Design Variable
 - Match voltage and frequency to required performance
- Minimize waste (or reduce switching capacitance)
 - Match computation and architecture
 - Preserve locality inherent in algorithm
 - Exploit signal statistics
 - Energy (performance) on demand

More easily accomplished in application-specific than programmable devices

Obviously misses the emerging importance of standby power ...

Panel ISLPED 96

(on the heels of the rowdy ISLPD 94 panel)

The holy grail of low-power systems

Design or Technology?

Jan M. Rabaey
U.C. Berkeley

Panel ISLPED

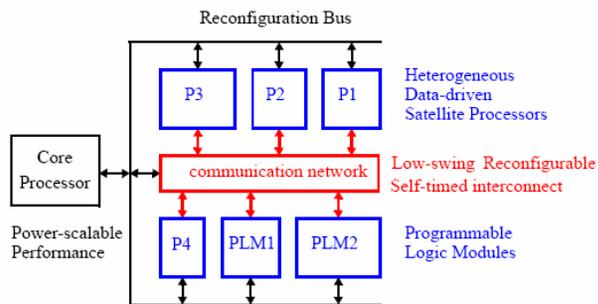
Observation:

Power is bound to become an issue
in virtually every system design

Some common deceptions ...

- Scaling the technology solves the power problem
Check every possible projection of microprocessor system power
- Advanced circuit design and better CAD tools have solved the problem
A drop on a boiling plate
- Power of electronics is not important anymore compared to display and electro-mechanical parts
Ratio of peripherals to core PD is continuously reducing

Energy-Efficient (Reconfigurable) Architectures



Combines: efficient no-frills computation, advanced circuit techniques, ease of use

A programmable voice coder @ 2mW

Panel ISLPED-96 - J. Rabaey

The state-of-the art in CAD

- Virtually all tools at circuit or gate level (the boiling plate ...) — a mixed response
- Only a single company offers architectural power analysis !
- Needs system level power analysis, prediction, budgetting, modeling and characterization (even at design conception level)
- Needs infrastructure for application-specific system on a chip design

Panel ISLPED-96 - J. Rabaey

Panel ISLPED 98

Past and Future BlockBusters in Low-Power Design



ISLPED 98
Evening Panel

Blockbuster events

- Reduction in supply voltage
- Architectural voltage scaling
- Low voltage-supply voltage processes
- Reduced voltage swing drivers
- Gated clocks
- On-chip PLLs
- Application specific architectural modifications
- Off-chip traffic minimization
- Optimal algorithms
- Power consumption simulation

Panel Composition

- Jan Rabaey, UC Berkeley — Moderator
- Bryan Ackland, Lucent — Signal processing, sensors
- Robert Brodersen, UC Berkeley — DSP, wireless
- Massoud Pedram, USC — CAD
- Christer Svensson, Linköping University — Digital circuits
- Bruce Wooley, Stanford University — Analog

Challenges

- Life after CMOS
- Dynamic voltage and frequency scaling
- Utilizing very low supply voltages
- Low power analog design
- Utilizing adiabatic techniques
- Low power tools in mainstream

>10 Years of Low-Power Design R&D

- Well on the road towards a structured low-power/energy design methodology!
 - From a grab-bag of techniques to modeling, simulation, estimation and synthesis techniques at different levels of the design hierarchy
 - Addressing both dynamic and static power
 - Still in need of some major advances, but the concepts are there

[EE Times: Design News](#)
Cadence 'kit' eases low-power IC design

[Richard Goering](#)
[EE Times](#)
(05/14/2007 9:00 AM EDT)

Hoping to make low-power IC design techniques more accessible, Cadence Design Systems this week is announcing a new low-power design kit that complements existing Cadence design tools. The kit is a "representative" design, example libraries.

The kit adds to the Low Power System Design Kit. That solution promises a complete design flow for low-power ICs, based on the Common Power System Design Kit. [Cadence](#) and is now managed by

[PRINT THIS STORY](#)
[SEND AS EMAIL](#)
[REPRINTS](#)

[EE Times: Design News](#)
System-level synthesis scheme homes in on low-power IC design

[Richard Goering](#)
[EE Times](#)
(05/14/2007 9:00 AM EDT)

Claiming a breakthrough approach to low-power IC design, ChipVision Design Systems this week will announce development of a patented low-power electronic system-level (ESL) synthesis technology. The company expects to roll out products based on the technology at the end of this year.

[PRINT THIS STORY](#)
[SEND AS EMAIL](#)
[REPRINTS](#)

[EE Times:](#)
Low-power IC design techniques may perturb the entire flow

[Richard Goering](#)
[EE Times](#)
(05/07/2007 9:00 AM EDT)

When NXP Semiconductors started to use advanced low-power IC design techniques, it was in for a surprise. "In some cases, we have experienced a twofold productivity drop for the implementation phase," said Hervé Menager, design and technology officer at NXP.

That's far from a unique experience. While EDA vendors have been fighting over two competing low-power specification [format](#) standards, a larger problem may have been obscured: Low-power techniques such as multivoltage design are so difficult that a

[PRINT THIS STORY](#)
[SEND AS EMAIL](#)
[REPRINTS](#)

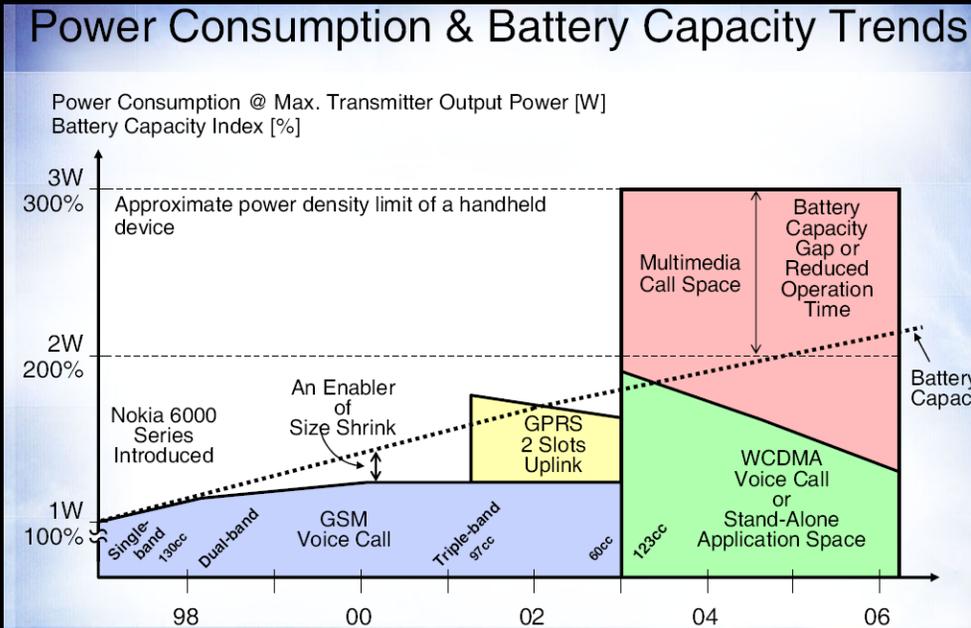
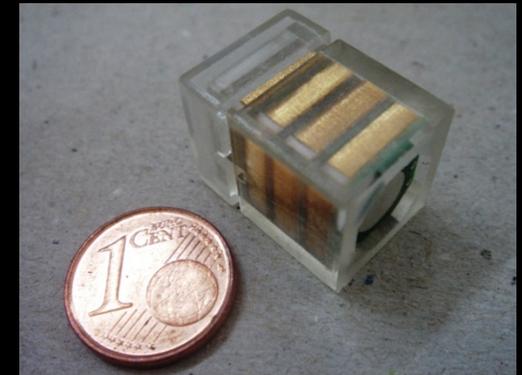
Power Now the Dominant Design Constraint



The Bad News

Google Data Center, The Dalles, Oregon

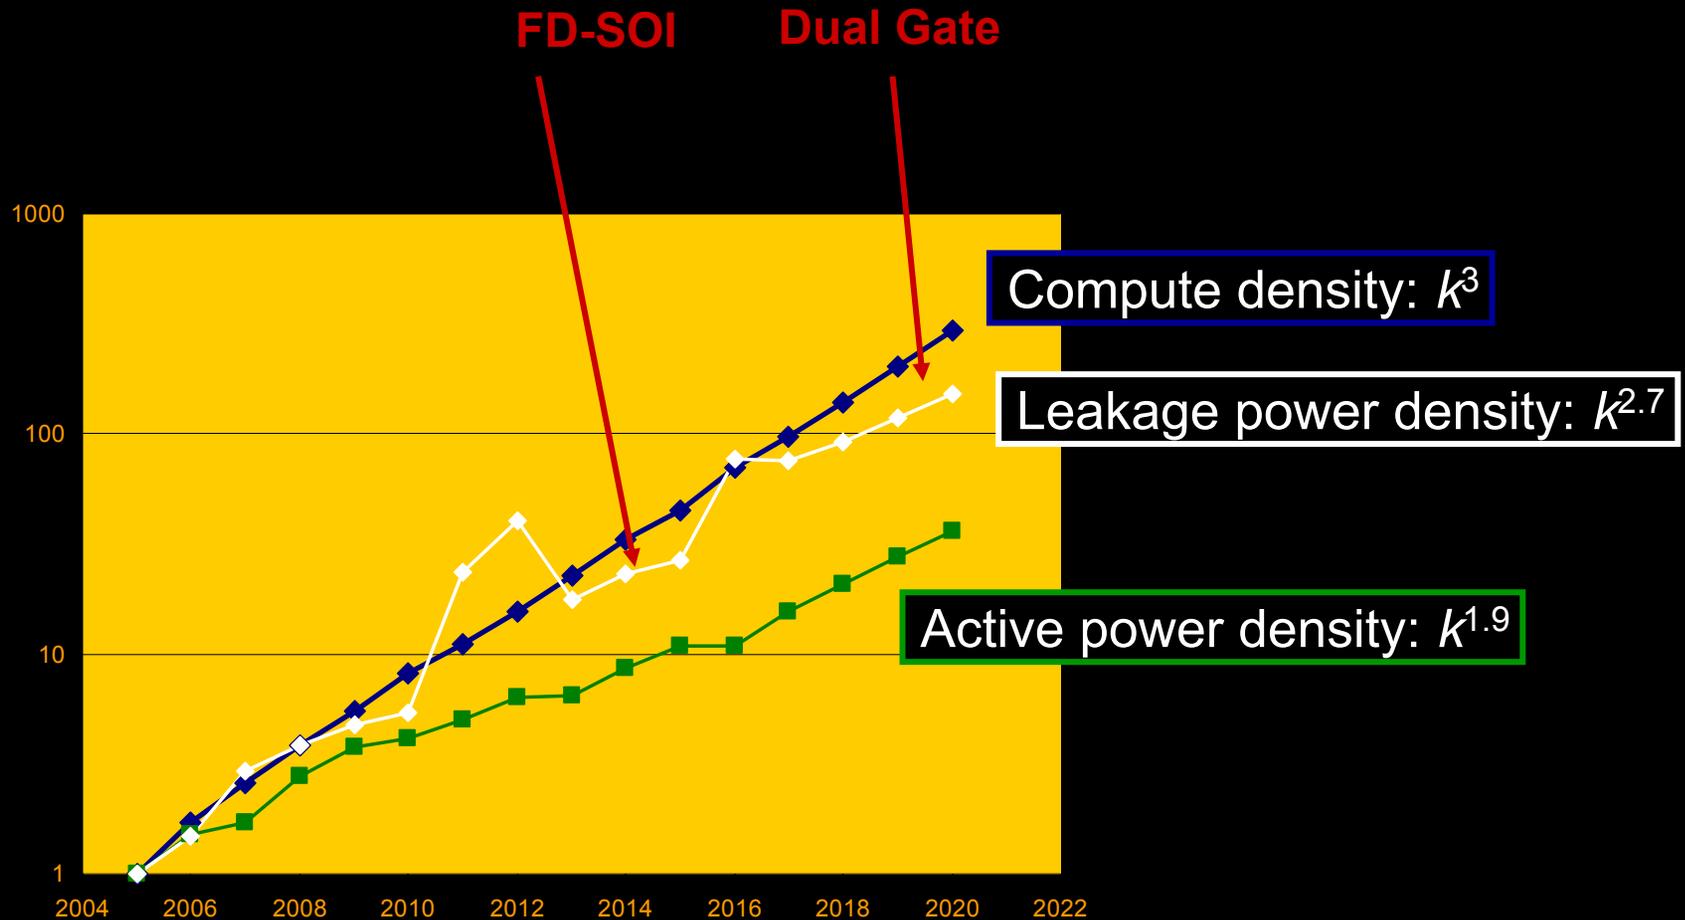
UCB PicoCube



Innovation necessary again

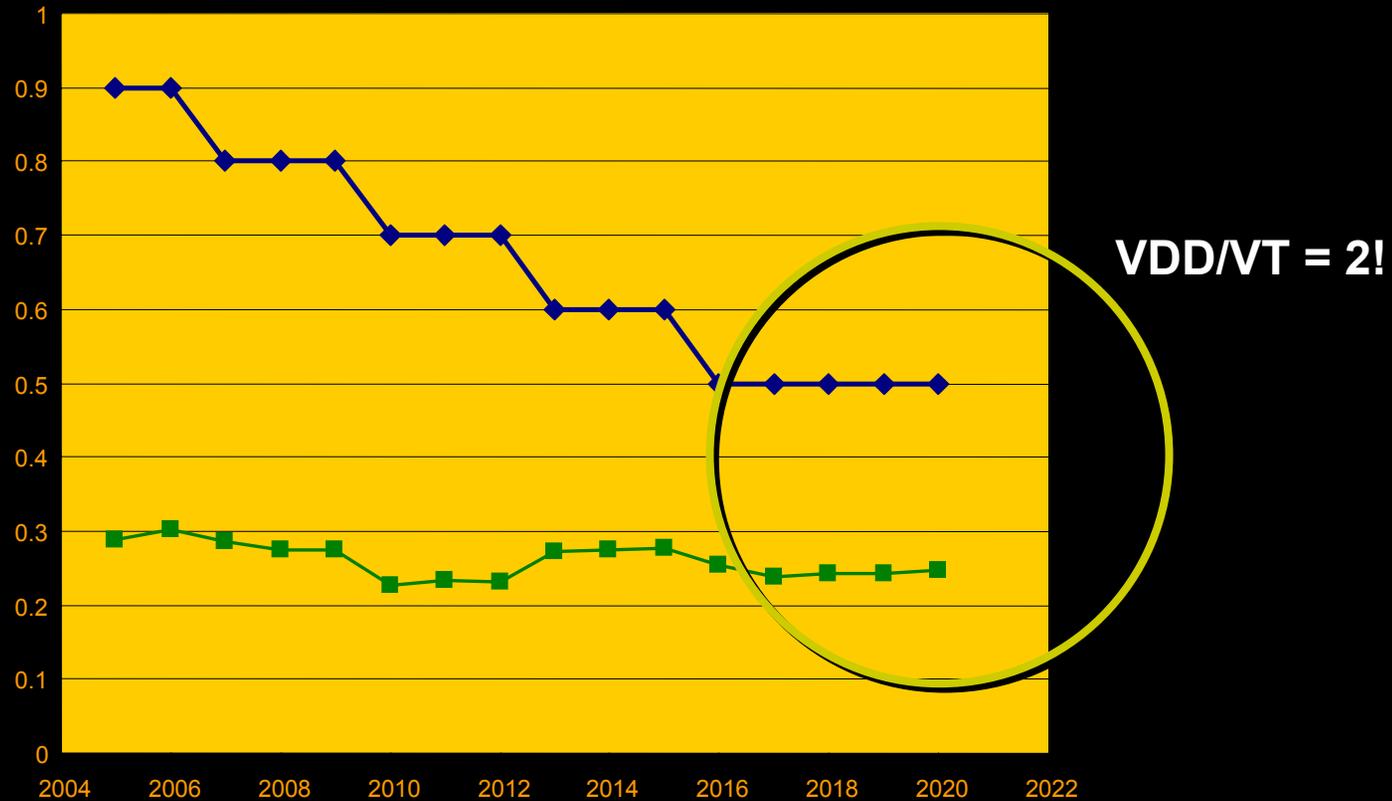
Y. Nuevo, ISSCC 04

Power and Energy Limiting Integration The Roadmap Perspective (2005)



Not looking good!
Technology innovations help, but impact limited.

Reducing Supply (and Threshold Voltages) an Essential Component



Optimistic scenario – some claims exist that VDD may get stuck around 1V

An Era of Power-Limited Technology Scaling

Technology innovations (will) offer some relief

- Devices that perform better at low voltage without leaking too much
- Example: FD-SOI, Dual-gate devices, Enhanced mobility transistors, MEMS-gate Devices

But also are adding major griève

- Impact of increasing process variations and various failure mechanisms more pronounced in low-power design regime.

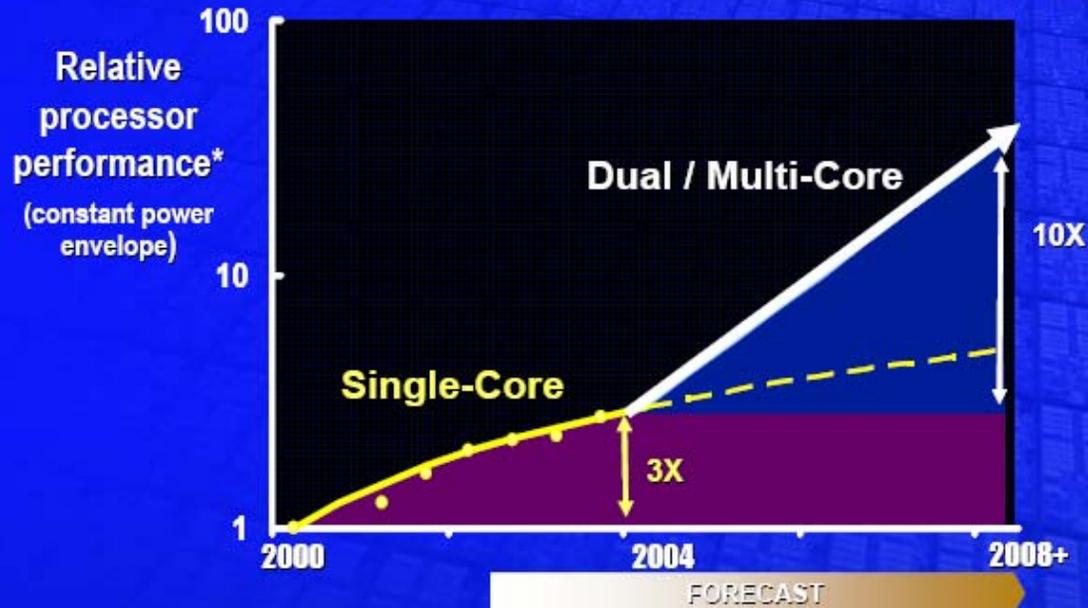
In dire need of new solutions if scaling is to continue

The Low Energy Roadmap

- **Concurrency Galore**
 - Many simple things are way better than one complex one
- **Always-Optimal Design**
 - Aware of operational, manufacture and environmental variations
- **Better-than-worst-case Design**
 - Go beyond the acceptable and recoup
- **Ultra-Low Voltages**
 - Exploring the boundaries
 - It might be easier than you think
- **Explore the Unknown**

Concurrency Galore

Performance and Power Efficiency Increase with Parallel Architecture



*Average of SPECint2000 and SPECfp2000 rates for Intel desktop processors vs initial Intel® Pentium® 4 Processor

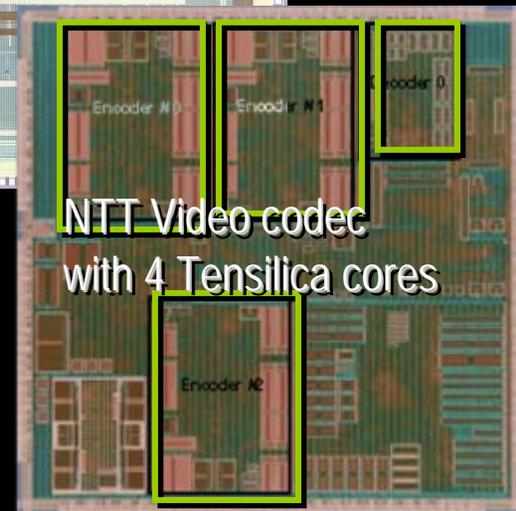
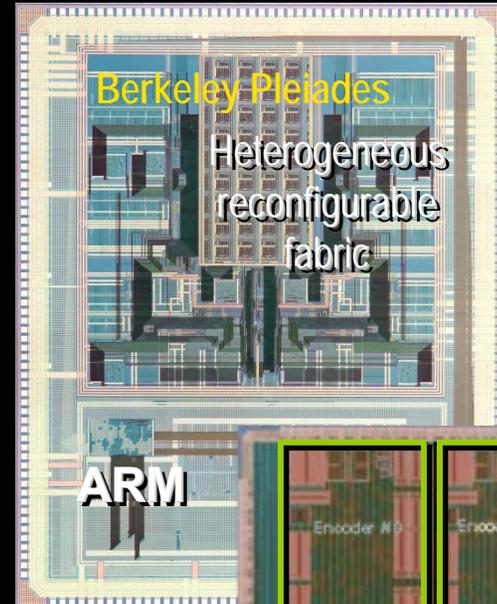
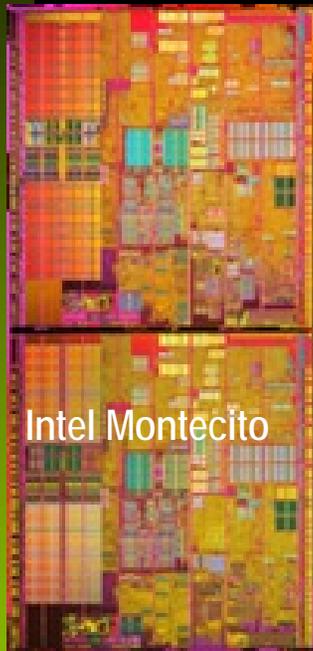
© 2005 IEEE International Solid-State Circuits Conference © 2005 IEEE

Source: Intel

Sunlin Chu, Intel, ISSCC05

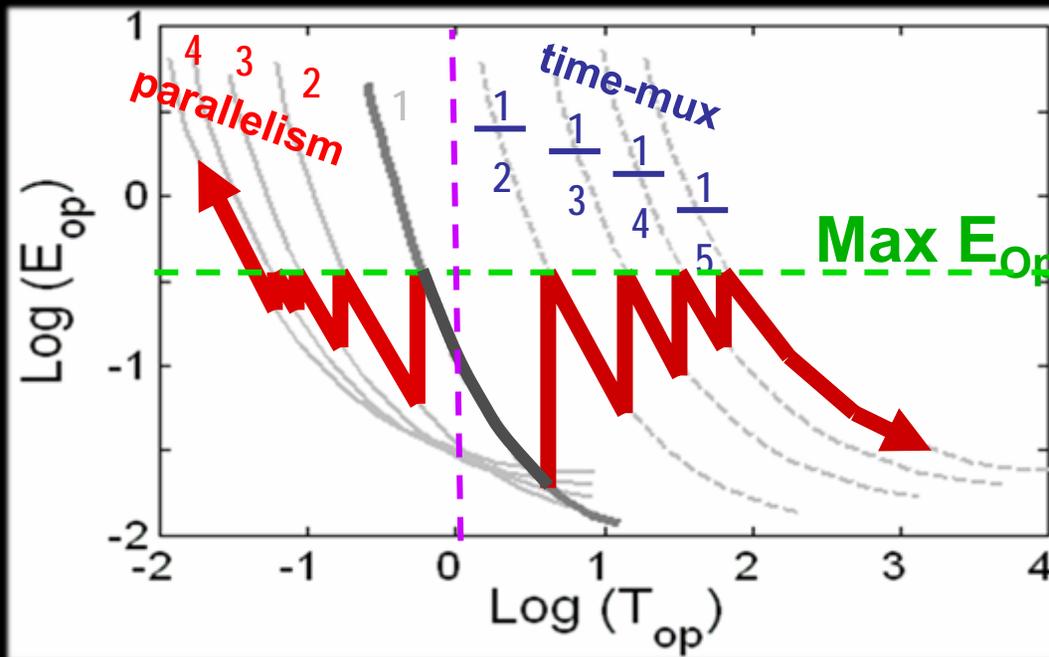
An obvious trend: more but simpler processors running at modest clock speeds and increased energy efficiency

Concurrent Multi-Many Core is Here to Stay



An obvious trend: more but simpler processors running at modest clock speeds and increased energy efficiency

The Underlying Story



Data for 64-b ALU
[Courtesy: Dejan Markovic]

LARGE



SMALL

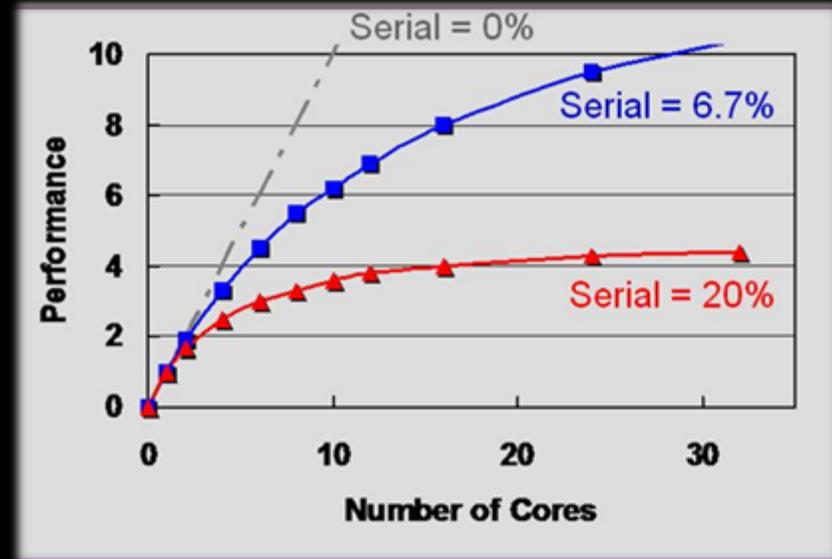
AREA

- For each level of performance, optimum amount of concurrency
- Concurrency provides higher performance for fixed energy/operation

OBSERVE: DOES NOT SCALE IN THE LONG TERM!

The Multi-Many Core Challenges

- In urgent need of software solution
 - Paradigm only works if sufficient concurrency is present!



- The architectural challenge

- What concurrent micro- and network architecture will prove to be ultimately viable: Multi-core versus reconfigurable, homogeneous versus heterogeneous, static versus dynamic routing
- What are the driving applications that are “massively parallel”
- Need exploration tools

Massive concurrency only makes sense if accompanied with simplification and voltage scaling, and overhead is bounded

The Multi-Core Reality

Paradigm only works if the concurrency is present and adequately exposed!

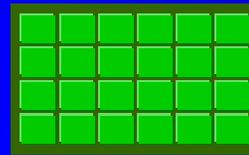
From Multi to Many...

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

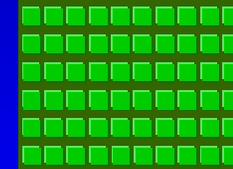
12 Cores



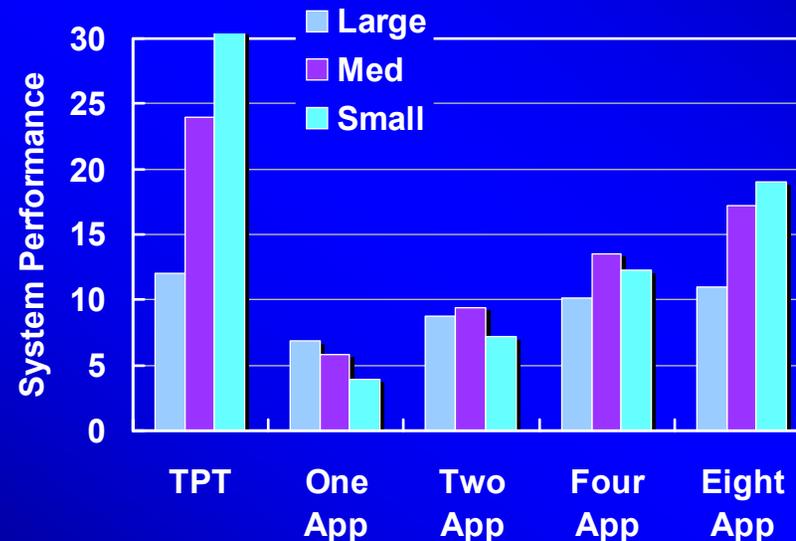
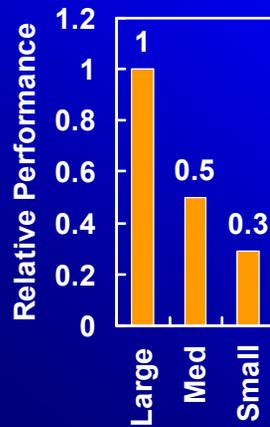
48 Cores



144 Cores



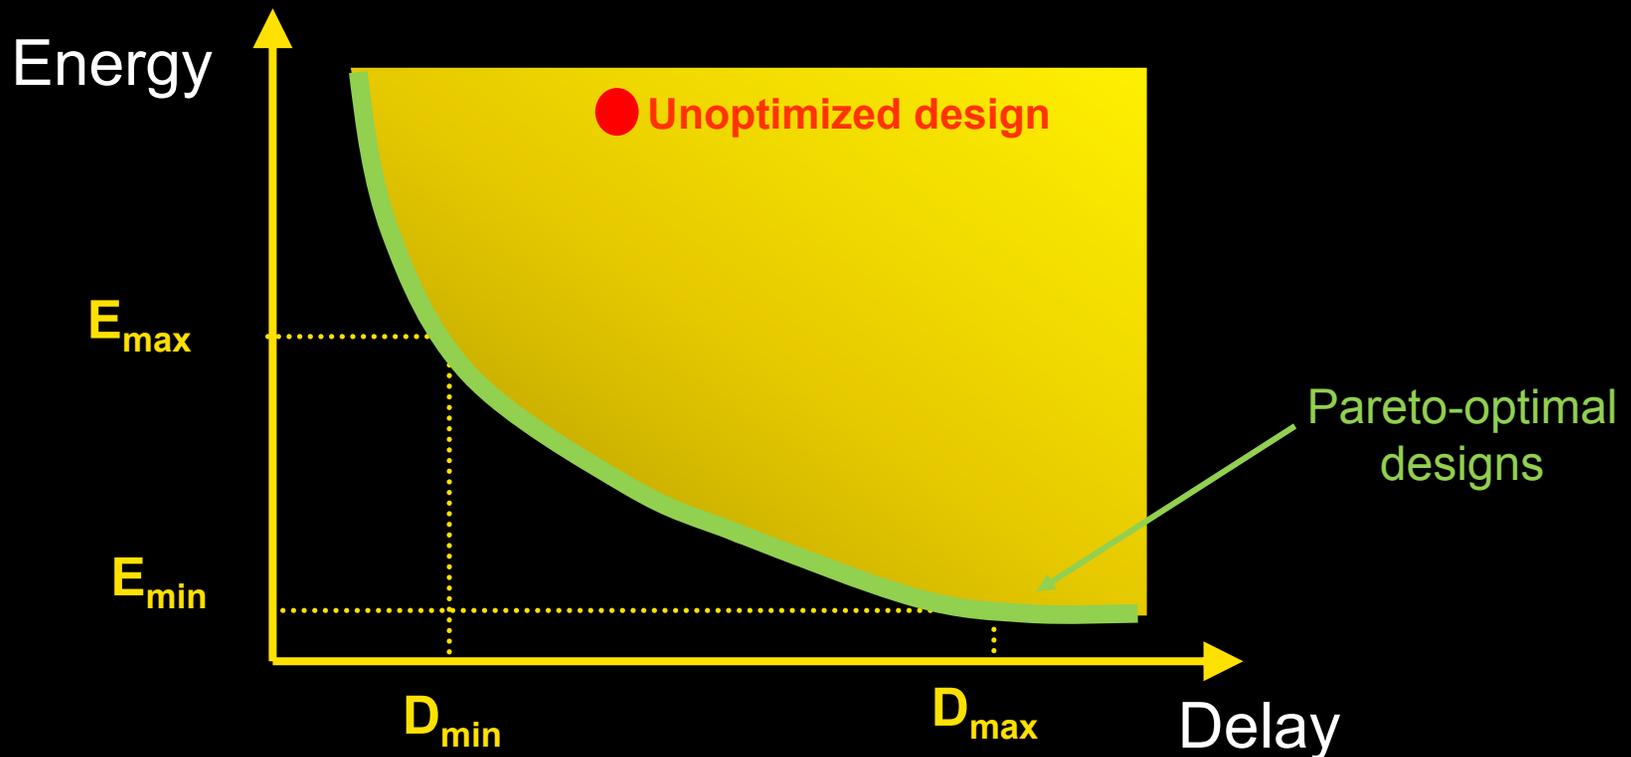
Single Core Performance



54

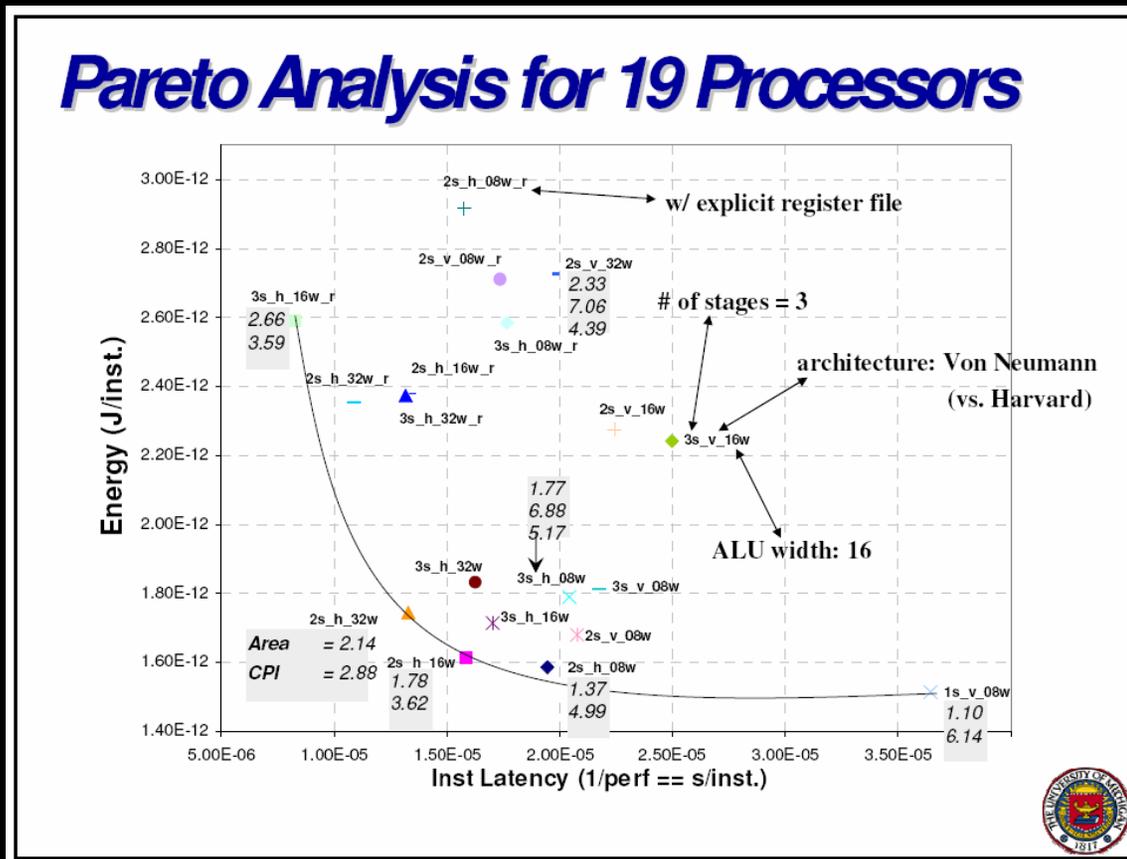
“Always-Optimal” Design

- For given function, activity and implementation instance, an **optimal operation point** exists in the energy-performance space



“Always-Optimal” Design

- For given function, activity and implementation instance, an **optimal operation point** exists in the energy-performance space



Simple is better from an energy perspective



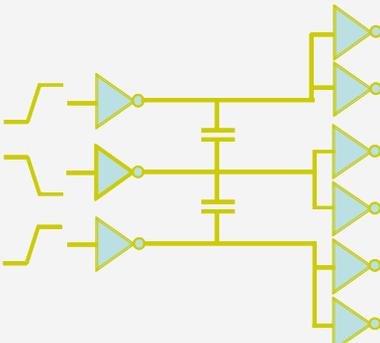
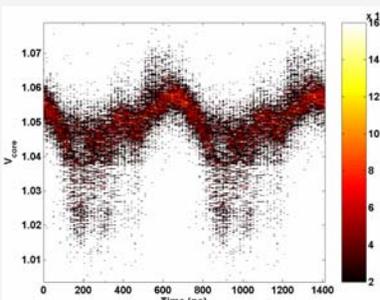
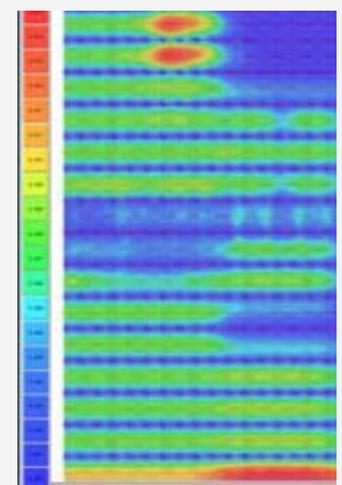
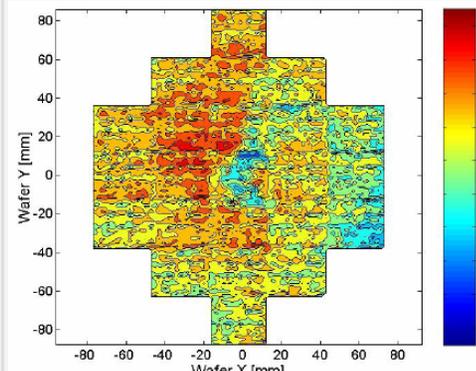
“Always-Optimal” Design

- For given function, activity and implementation instance, an **optimal operation point** exists in the energy-performance space
- Time of optimization depends upon **activity profile**
- Different optimizations apply to active and static power

	Fixed Activity	Variable Activity	No Activity - Standby
Active			
Static	Design time	Run time	Sleep

Energy-optimized systems must operate at optimal settings at every activity level → **run-time optimization!**

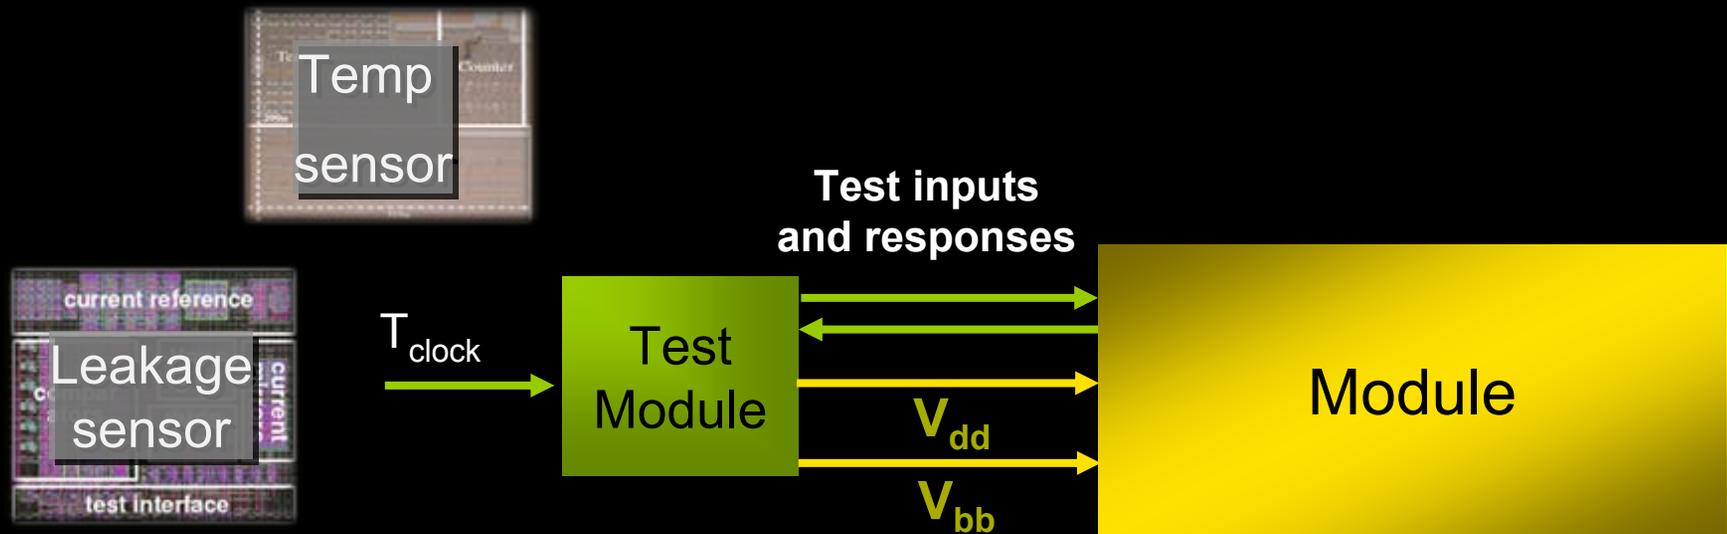
Adding Temporal and Spatial Variations

Signal Coupling	Supply/Package Noise	Temperature — Modal Operation	Manufacturing — Wear-out
10^{-10}-10^{-8}	10^{-7}-10^{-5}	10^{-4}-10^{-2}	10^5-10^7
			

Always-Optimal Systems

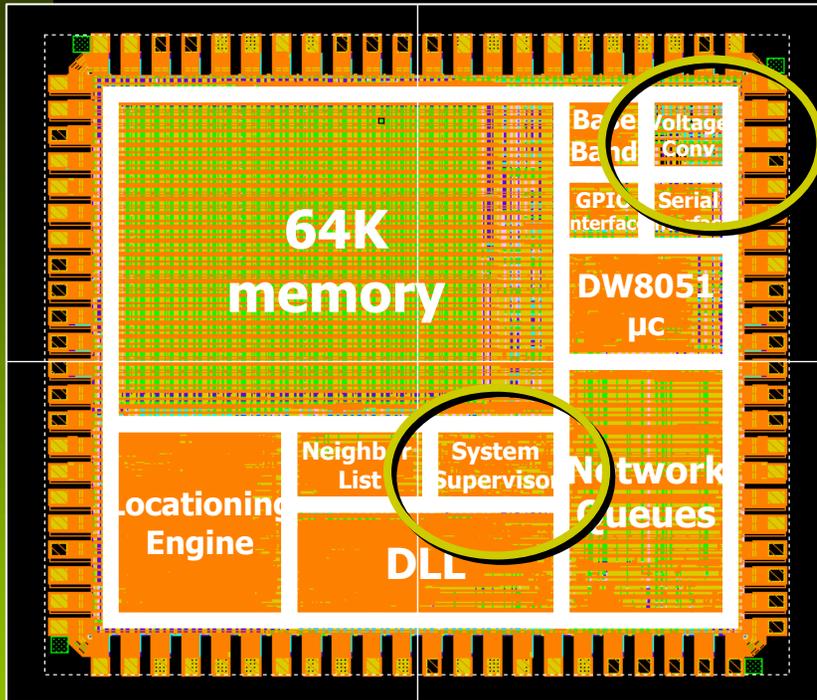
System modules are adaptively biased to adjust to operating, manufacturing and environmental conditions

- Parameters to be measured: temperature, delay, leakage
- Parameters to be controlled: V_{DD} , V_{TH} (or V_{BB})



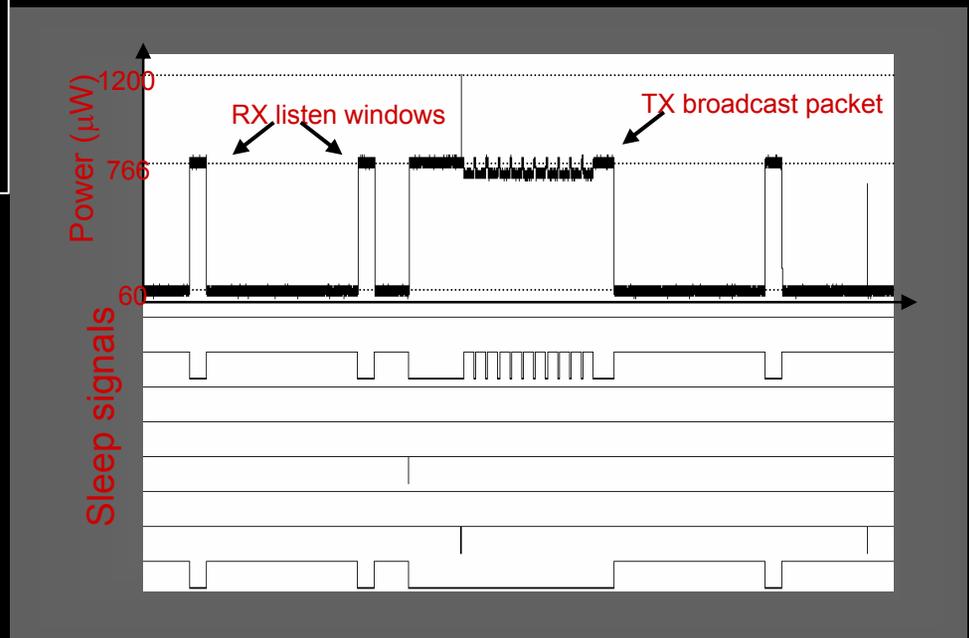
- Maximum power saving under technology and manufacturing limits
- Inherently improves the robustness of design timing
- Minimum design overhead required over the traditional design methodology

Extrapolates the Power Management Idea



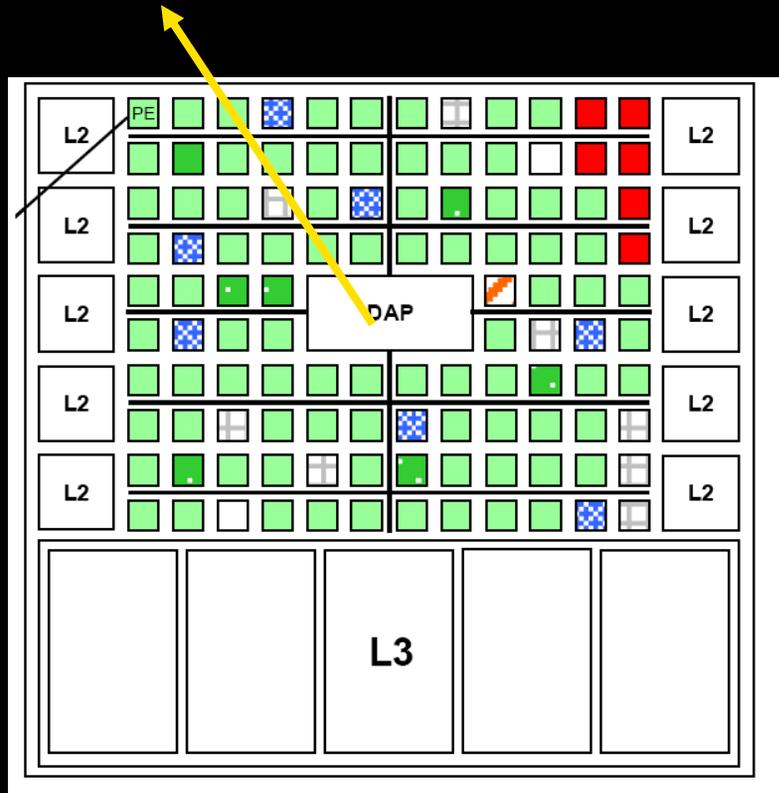
System supervisor evaluates and predicts activity and schedules voltage modes based on computational needs as well as measured parameters

Integrated Processor
for Sensor Networks,
M. Sheets, UCB



ElasticC – An “Always Optimal” IC

Diagnostic Adaptivity Processor



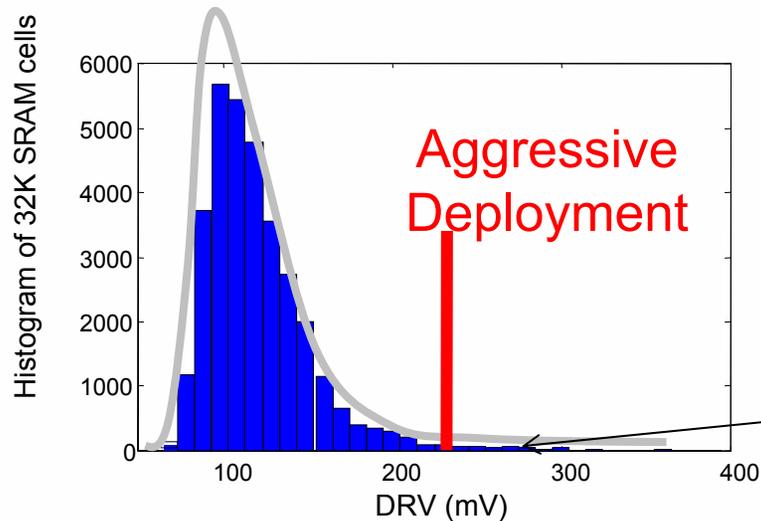
Multi-Core Architecture for Adaptability

- Monitor Temperature, Power, Reliability Degradation and Performance
- Provide real-time information to thread scheduling facilities
- Maintain system targets under varying stress conditions and usage profiles

- Needs architecture level perspective
- Challenges traditional test and verification flows

Better-than-worst-case design

- Also known as “**Aggressive Deployment (AD)**”
- Observation:
 - Current design targets worst case conditions, which are rarely encountered in actual operation
- Remedy:
 - Operate circuits at lower voltages level than allowed by worst case and deal with the occasional errors in other ways



Example:

Operate memory at voltages lower than allowed by worst case, and deal with the occasional errors through error-correction

Distribution ensures that error-rate is low

Better-than-worst-case Design — Components

Every aggressive deployment scheme must include the following components

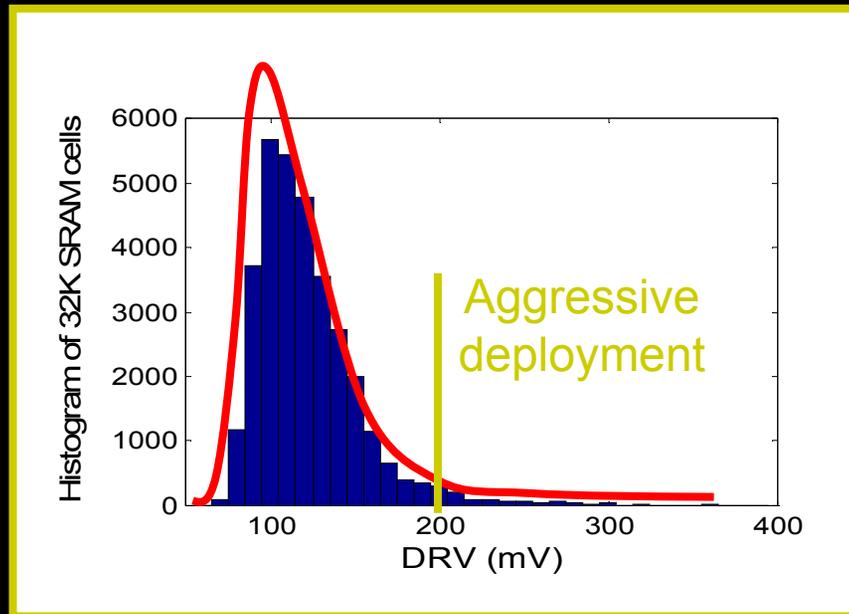
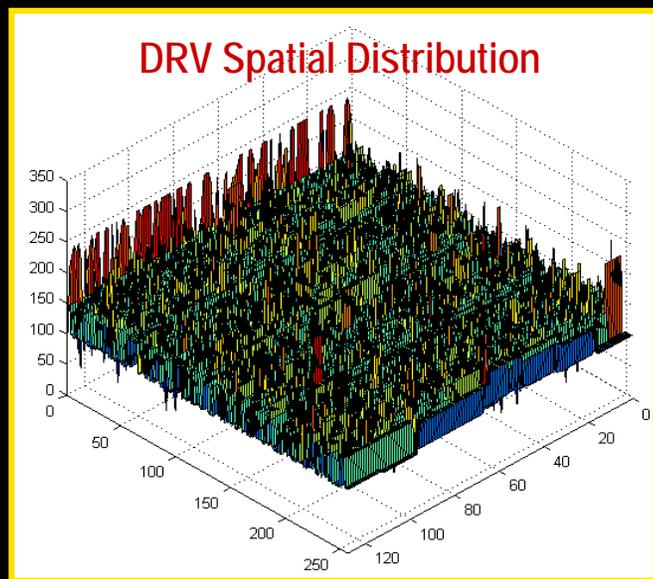
- **Voltage-setting Mechanism**
 - Distribution profile learned through simulation or dynamic learning
- **Error Detection**
 - Simple and energy-efficient detection is crucial for aggressive deployment to be effective
- **Error Correction**
 - Since errors are rare, its overhead is only of secondary importance

Concept can be employed at many layers of the abstraction chain (circuit, architecture, system)

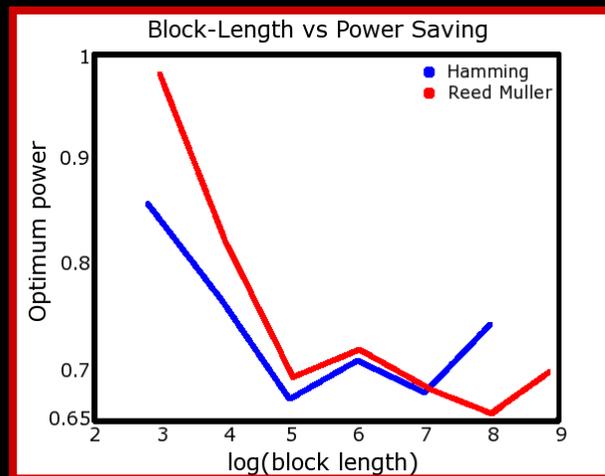
Aggressive Deployment

Operate circuits at voltages that are lower than worst case and deal with the occasional errors in other ways

Example: SRAM memory



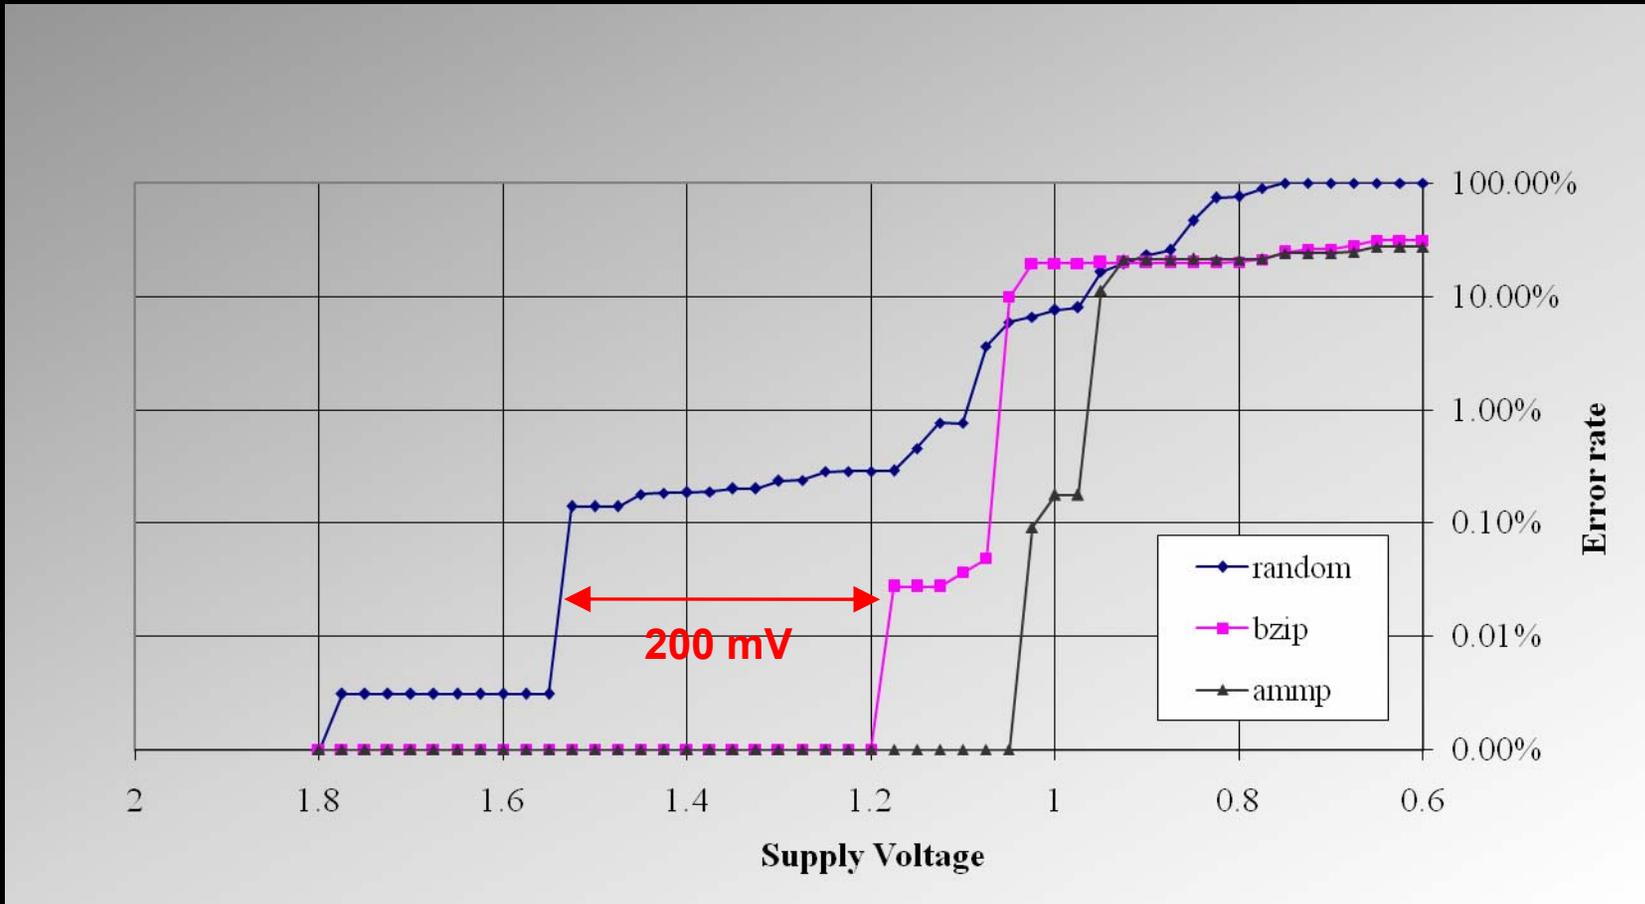
Hamming [31, 26, 3] : 33% power savings
Reed-Muller [256, 219, 8]: 35% savings



Source: Huifang Qin, ISQED 2004

Error Rate versus Supply Voltage

Example: Kogge-Stone adder (870 MHz) (SPICE Simulations) with realistic input patterns

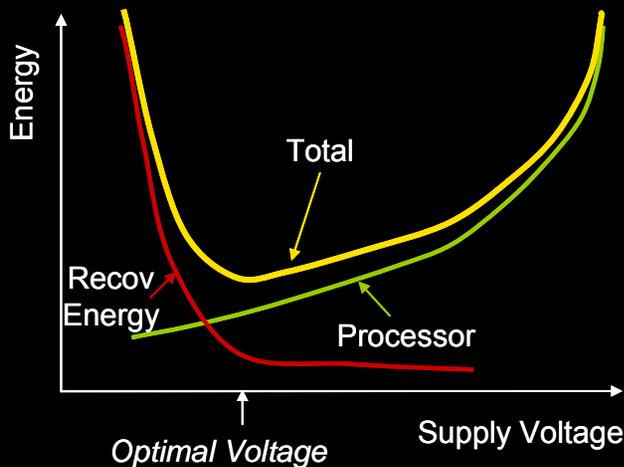
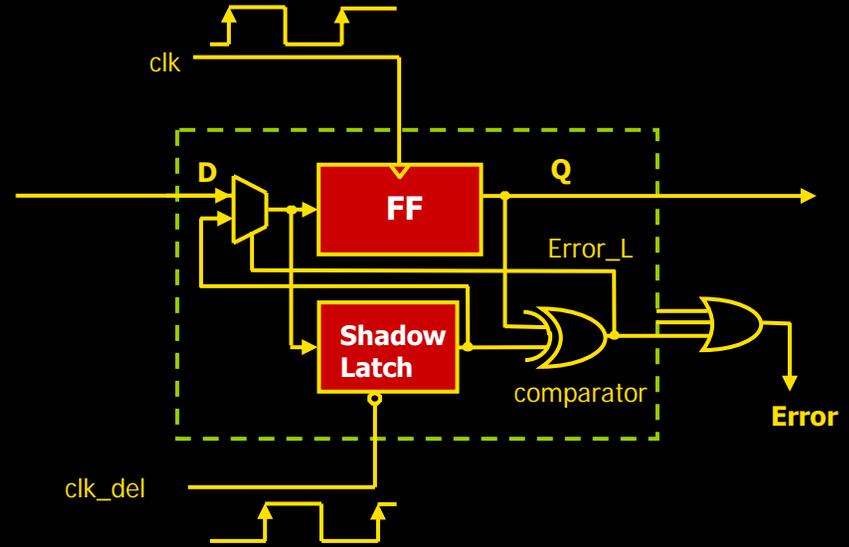


Better-than-worst-case Design

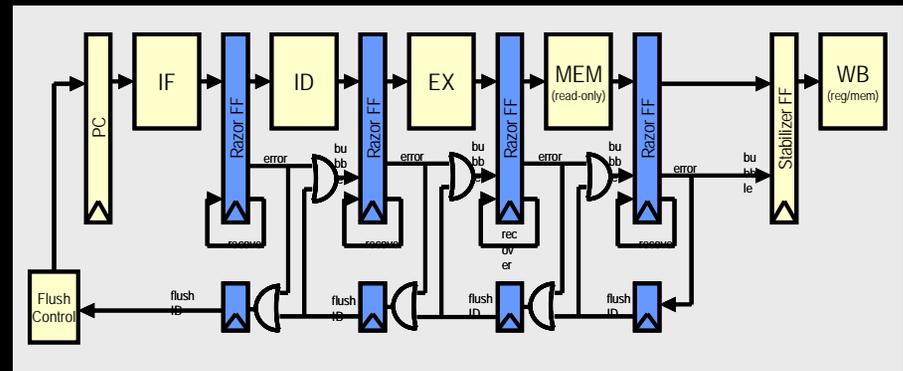
Scale voltage more than is allowable and deal with the consequences

Example: “Razor”

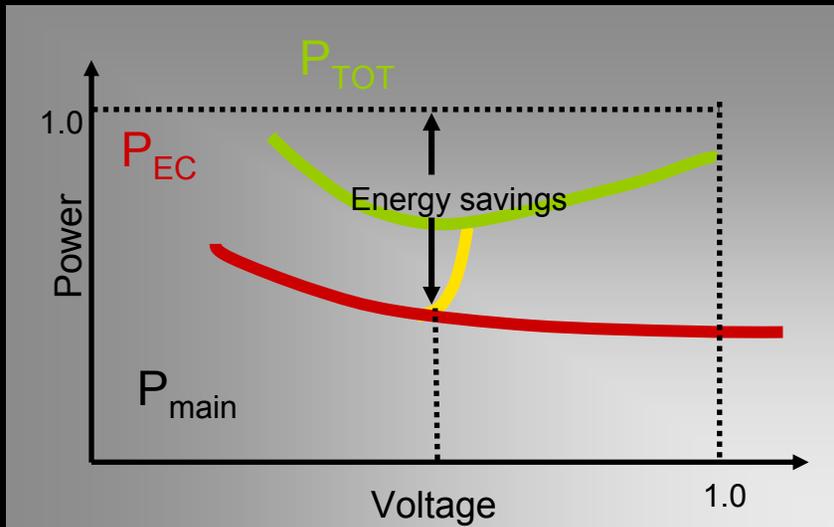
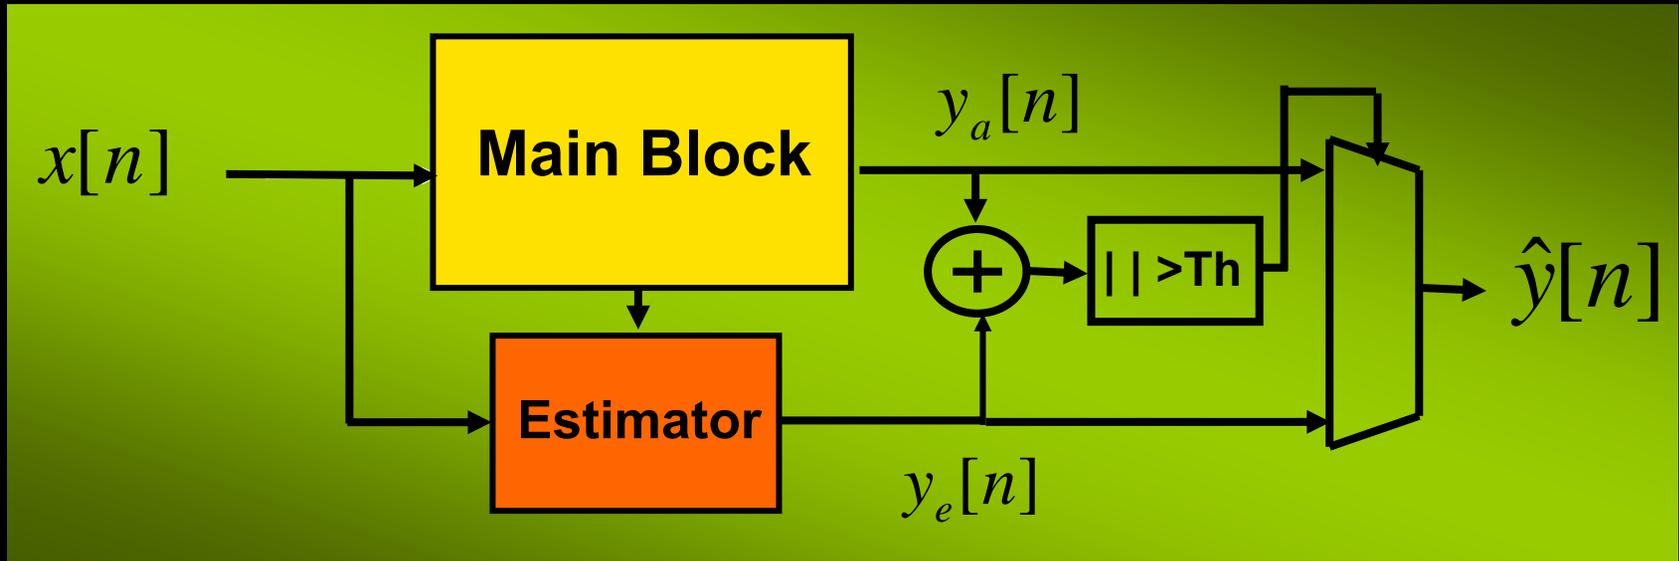
A “pseudo-synchronous” approach to address process variations and power minimization with minimal overhead by combining circuit and architectural techniques



“razored pipeline”



“Aggressive” Deployment At the Algorithm Level



Voltage **overscale** Main Block.
Correct errors using Estimator.
Power savings $\geq 3X!$

Leveraging resiliency to increase value

error-free

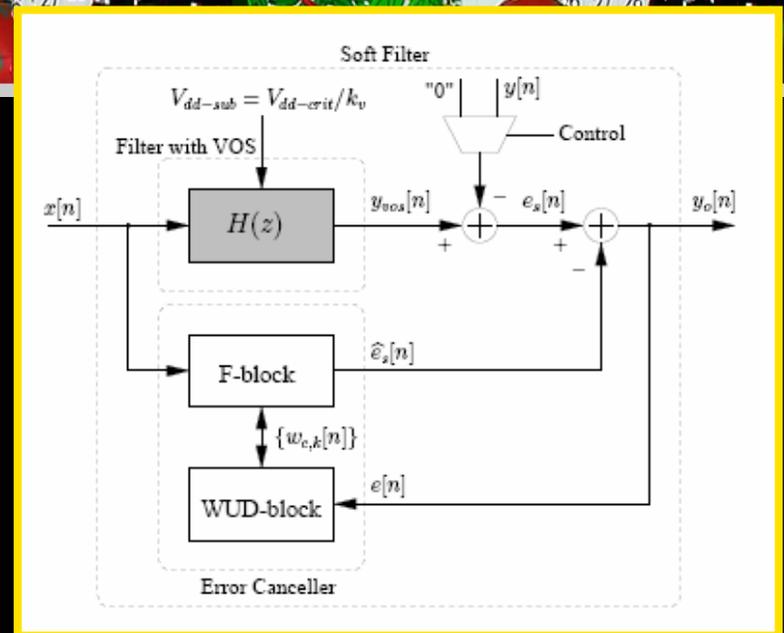
with errors

error-corrected

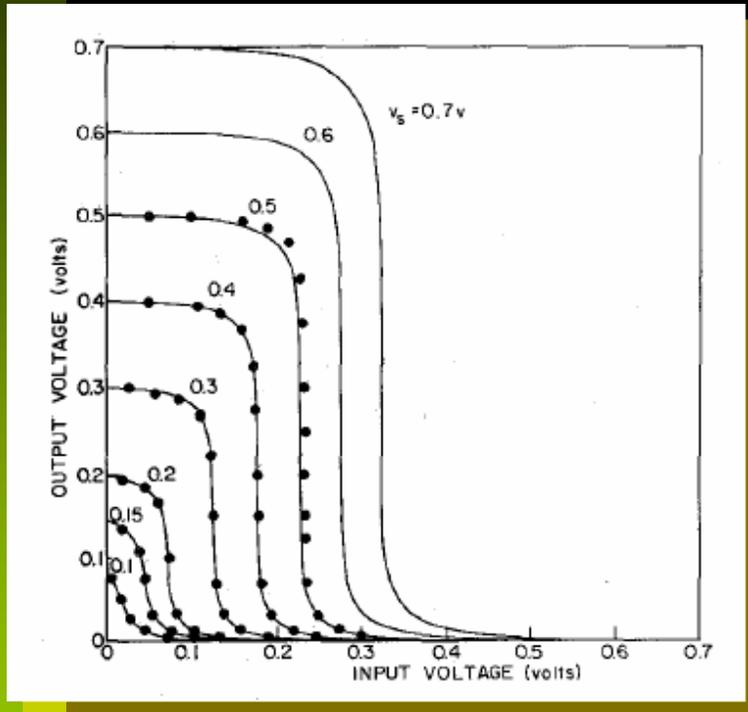


Low power motion estimation architecture using Algorithmic Noise Tolerance (Shanbhag, UIUC)

Up to 71% energy reduction demonstrated



Ultra-Low Voltage Design – Aggressive Deployment to the Extreme



[Swanson, Meindl (1972, 2000)]

**There is room at the
bottom**

Minimum operational voltage (ideal MOSFET):

$$V_{dd}(\min) \cong 2 (\ln 2)kT/q = 1.38kT/q = 0.036 \text{ V at } 300^\circ\text{K}.$$

Minimum energy/operation = $kT\ln(2)$

5 orders of
magnitude below
current practice
(90 nm at 1V)



[Von Neumann (1966)]

Equivalence between Communication and Computation

Shannon's theorem on maximum capacity of communication channel

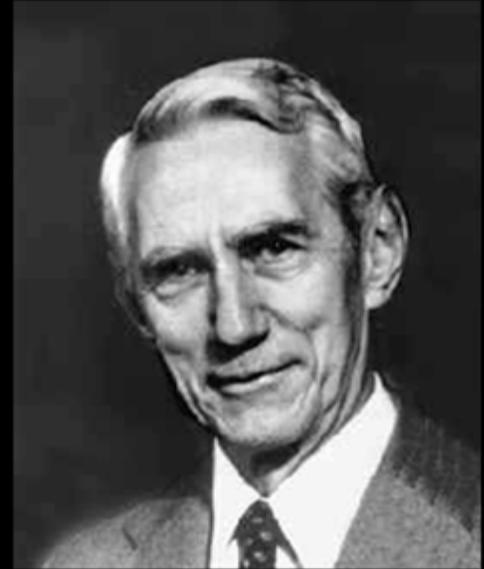
$$C \leq B \log_2 \left(1 + \frac{P_s}{kTB} \right)$$

$$E_{bit} = P_s / C$$

C : capacity in bits/sec

B : bandwidth

P_s : average signal power



Claude Shannon

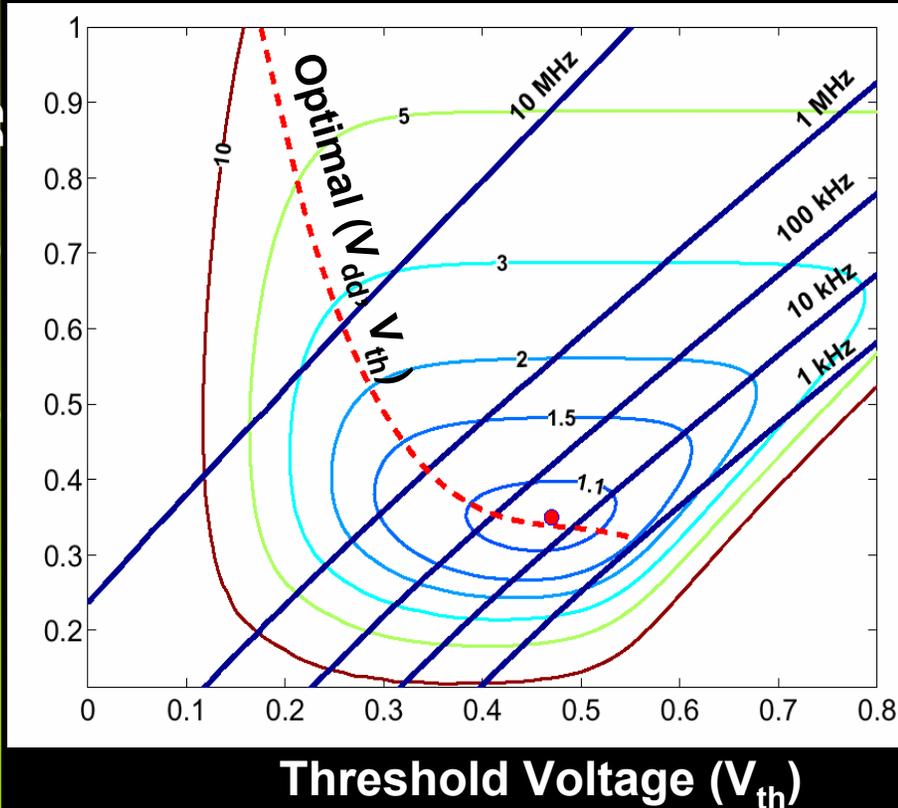
$$E_{bit} (\text{min}) = E_{bit} (C / B \rightarrow 0) = kT \ln(2)$$

Valid for an "infinitely long" bit transition ($C/B \rightarrow 0$)

Equals $4 \cdot 10^{-21}$ J/bit at room temperature

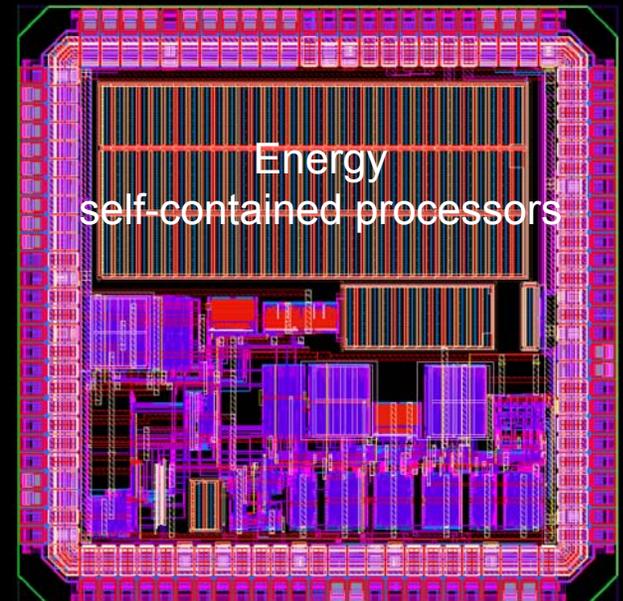
Sub-Threshold Leads to Minimum Energy/Operation

Supply Voltage (V_{DD})



Energy-Aware FFT Processor
[Chang, Chandrakasan, 2004]

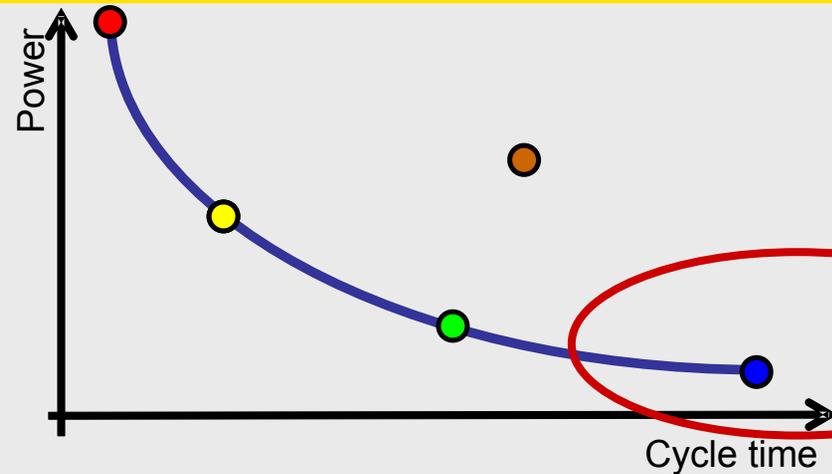
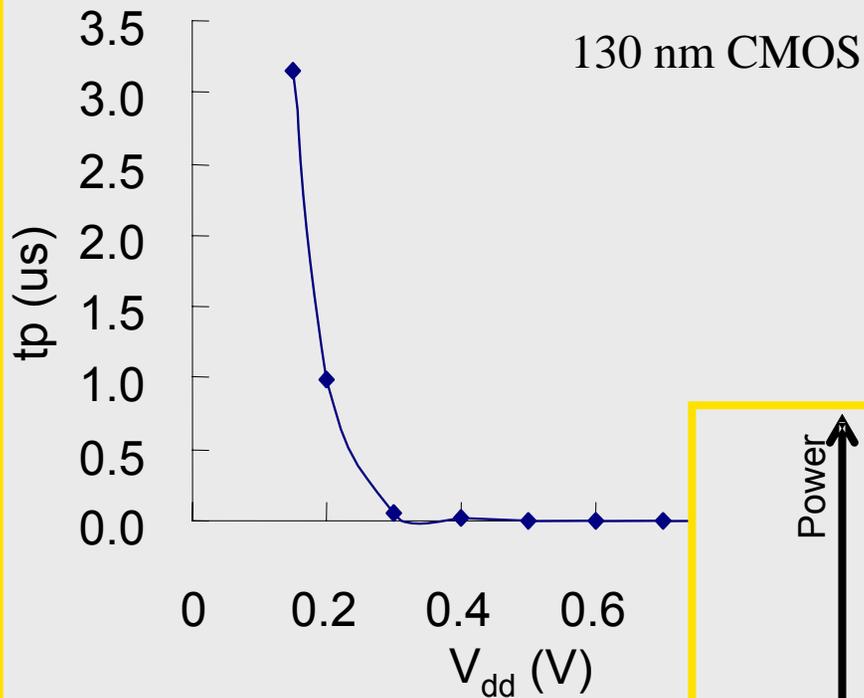
But ... At a huge cost in
performance



Subliminal processor
[Blaauw, 2006]
3 pJ/inst @ 350 mV

Is Sub-threshold The Way to Go?

- Achieves lowest possible energy dissipation
- But ... at a dramatic cost in performance



— OPTIMAL POWER – PERFORMANCE TRADEOFF CURVE

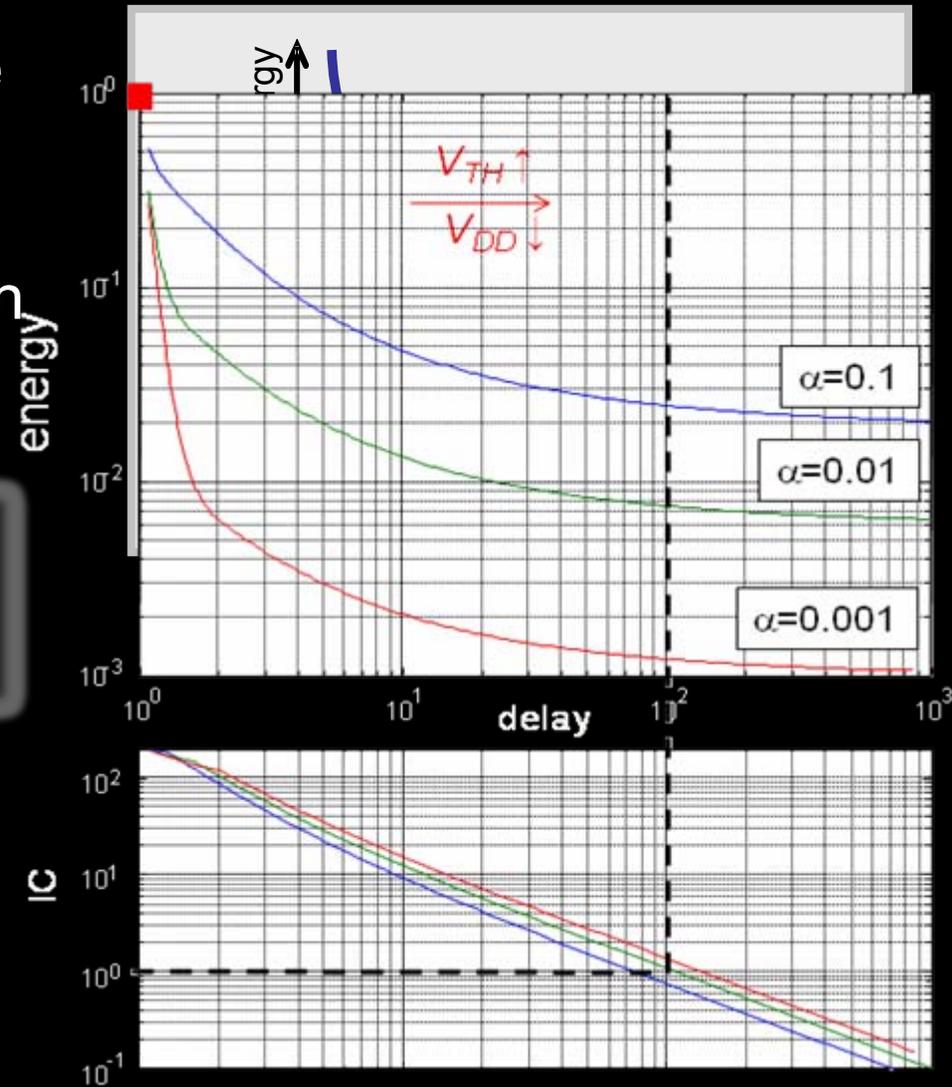
Backing Off a Bit

- Operating slightly above the threshold voltage improves performance dramatically while having small impact on energy

The Challenge: Modeling and Design in the Weak and Moderate Inversion Region

It is easier than you think!!

Example: optimization of adder over full design space (V_{DD} , V_T , W) using EKV model

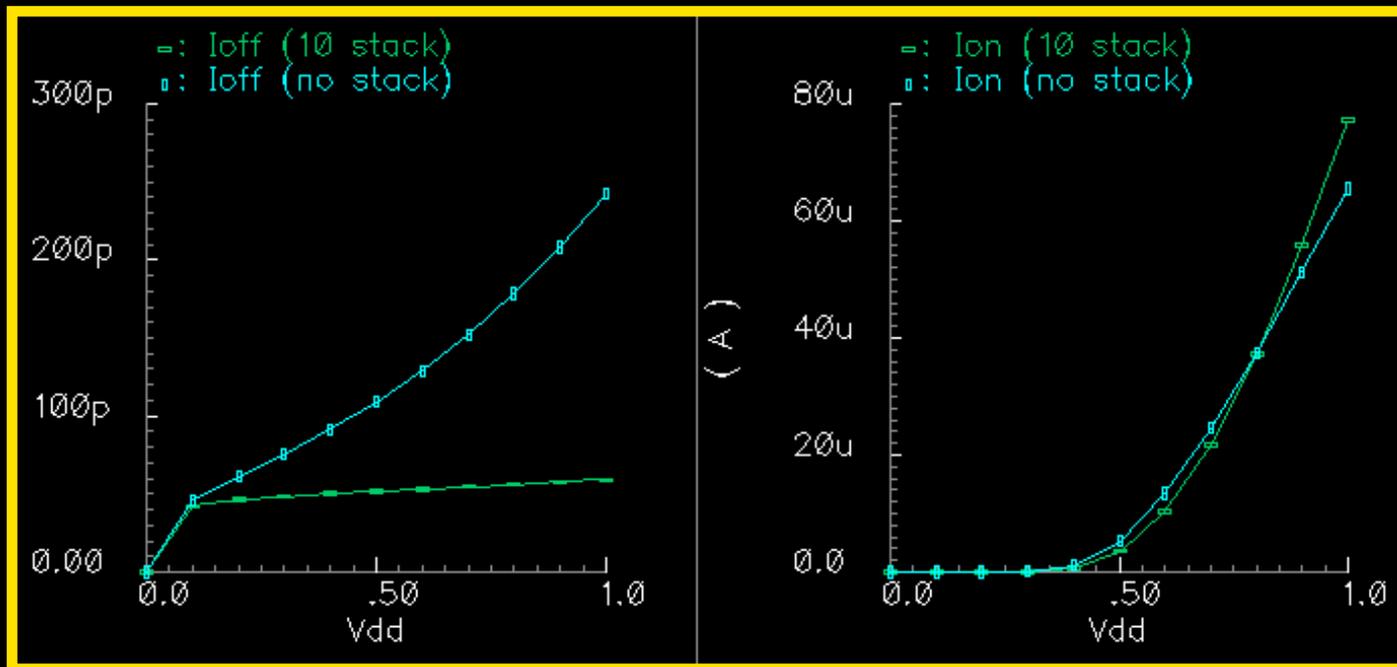


Need to Scale Thresholds as Well!

But need to manage leakage.

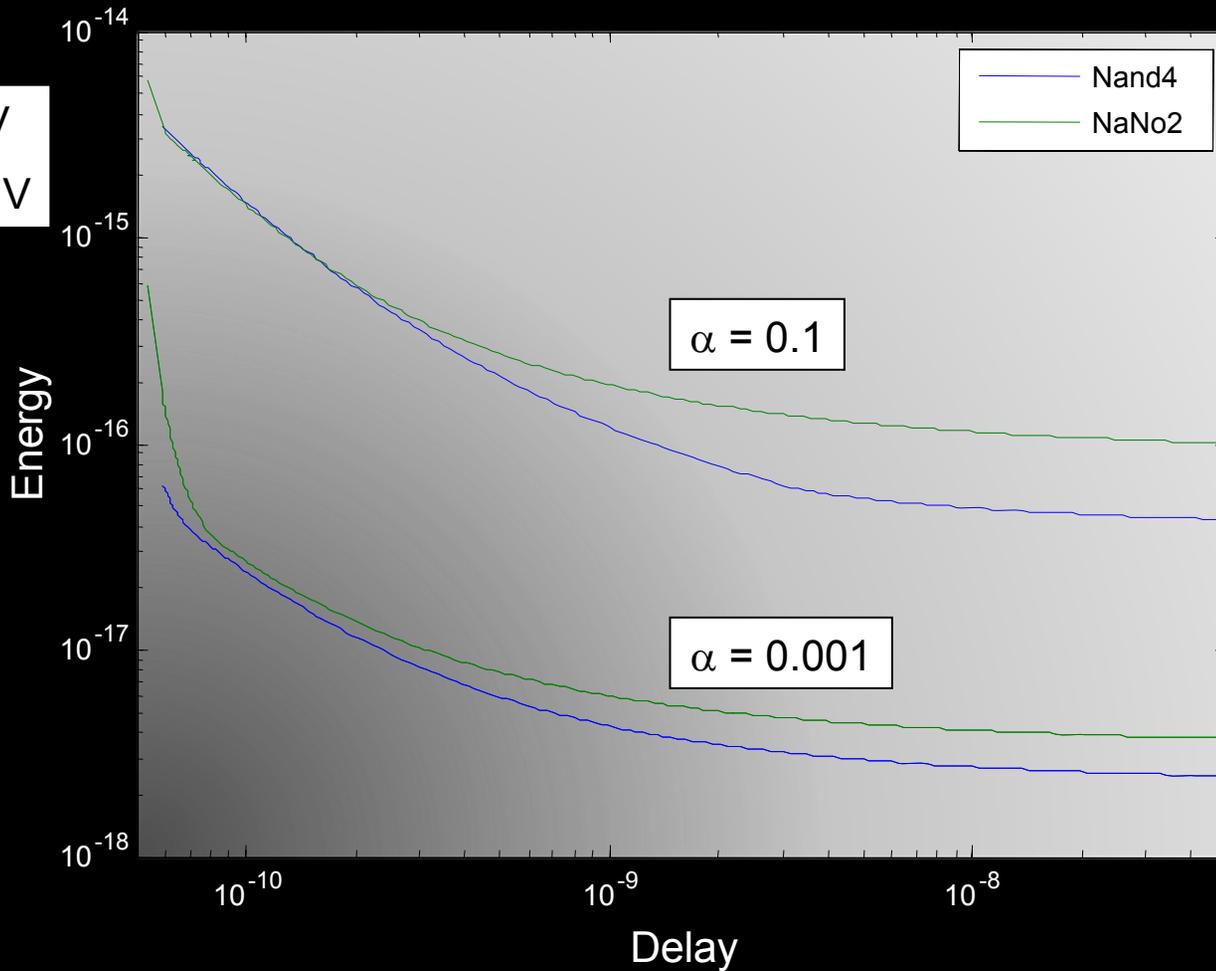
One option: Stacked transistors

- Ion/Ioff increases with increasing stack height (leakage suppression)
- More robust to correlated (tune or adapt) and random variations (self-cancel)
- Decreased short channel effect



Complex versus Simple Gates

$V_{dd} = 1V$
 $V_{TH} = 0.1V$



$V_{dd} = 0.14V$
 $V_{TH} = 0.25V$

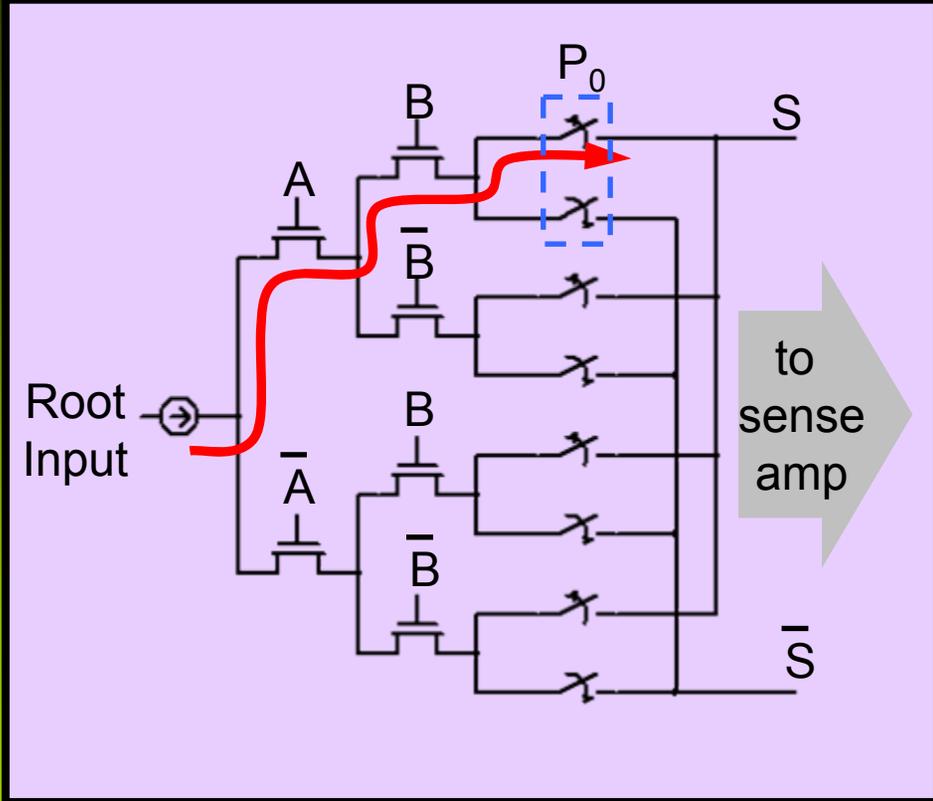
$V_{dd} = 0.1V$
 $V_{TH} = 0.22V$

$V_{dd} = 0.34V$
 $V_{TH} = 0.43V$

$V_{dd} = 0.29V$
 $V_{TH} = 0.38V$

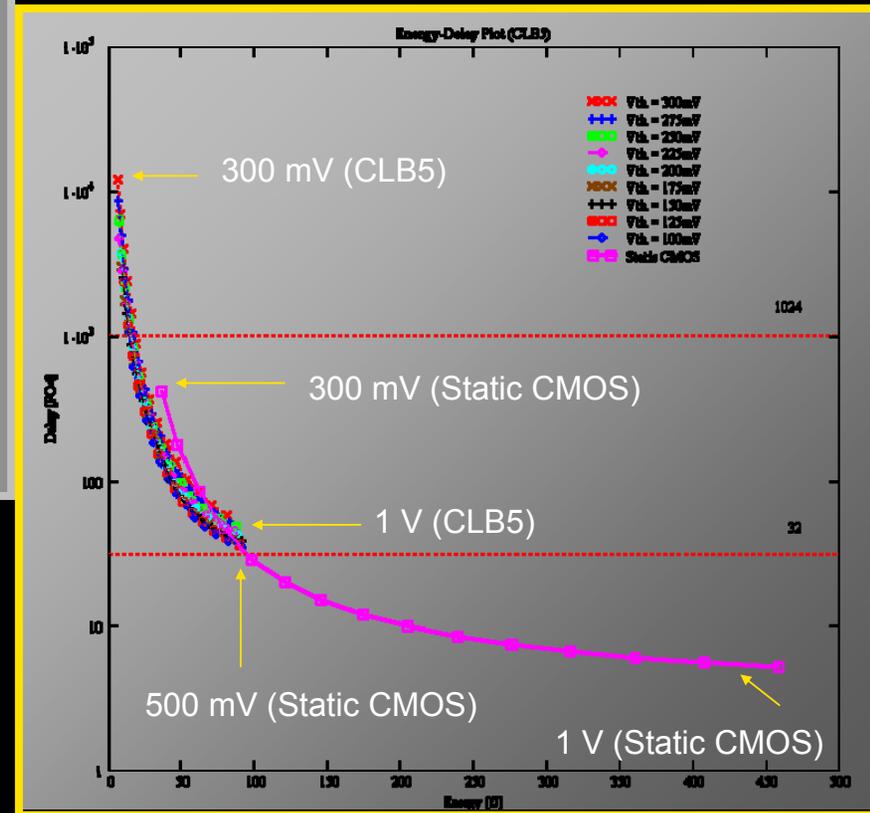
Complex Gates

Reducing thresholds while containing leakage



Example: pass-transistor logic

- Current-steering
- Regular
- Balanced delay
- Programmable



Exploring the Unknown – Alternative Computational Models

Humans



- 10-15% of terrestrial animal biomass
- 10^9 Neurons/"node"
- Since 10^5 years ago

The Yellow Brick Road of Ultra Low-Power Design

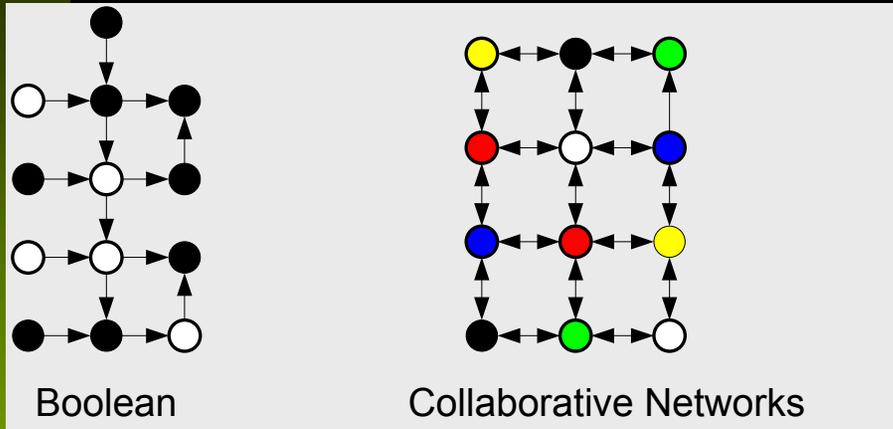
Ants



- 10-15% of terrestrial animal biomass
- 10^5 Neurons/"node"
- Since 10^8 years ago

Easier to make ants than humans
"Small, simple, swarm"

Example: Collaborative Networks



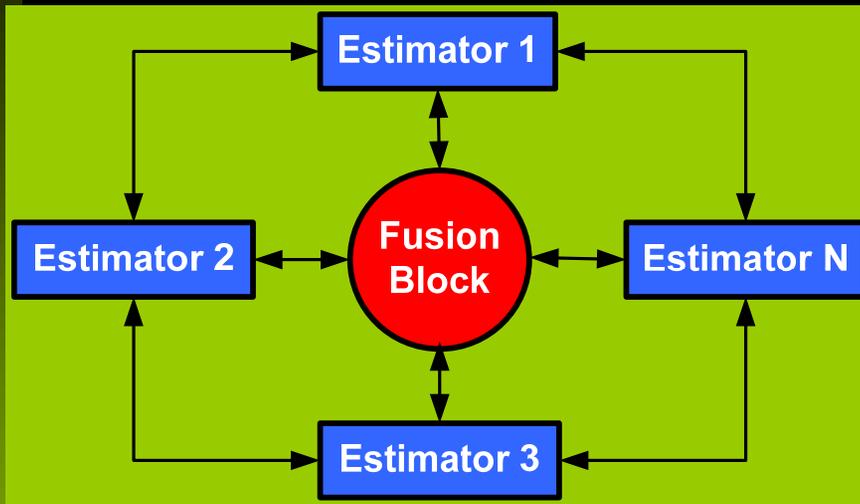
Metcalfe's Law
to the rescue of
Moore's Law!



Bio-inspired

- Networks are intrinsically robust → exploit it!
- Massive ensemble of cheap, unreliable components
- Network Properties:
 - Local information exchange → global resiliency
 - Randomized topology & functionality → fits nano properties
 - Distributed nature → lacks an “Achilles heel”

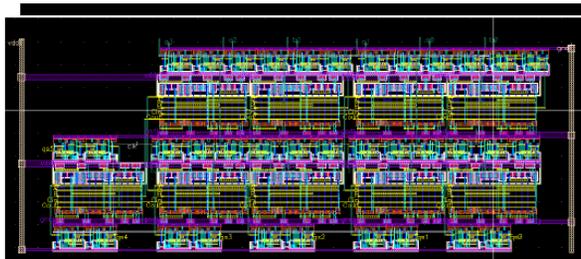
Example: “Sensor Networks on a Chip”



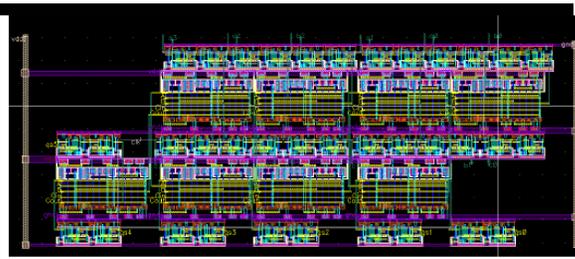
Use “large” number of very simple unreliable components

Estimators need to be independent for this scheme to be effective

Architecture Diversity



Dual Ripple Carry



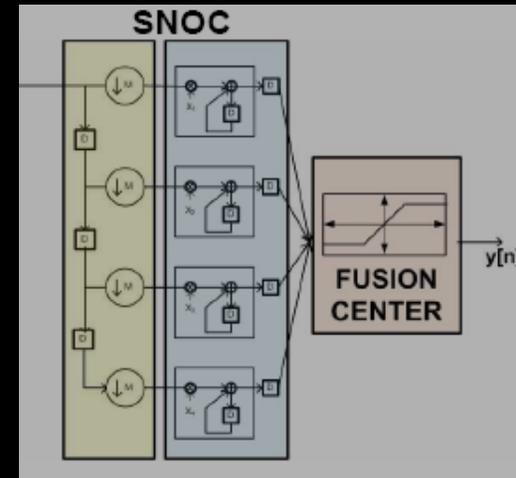
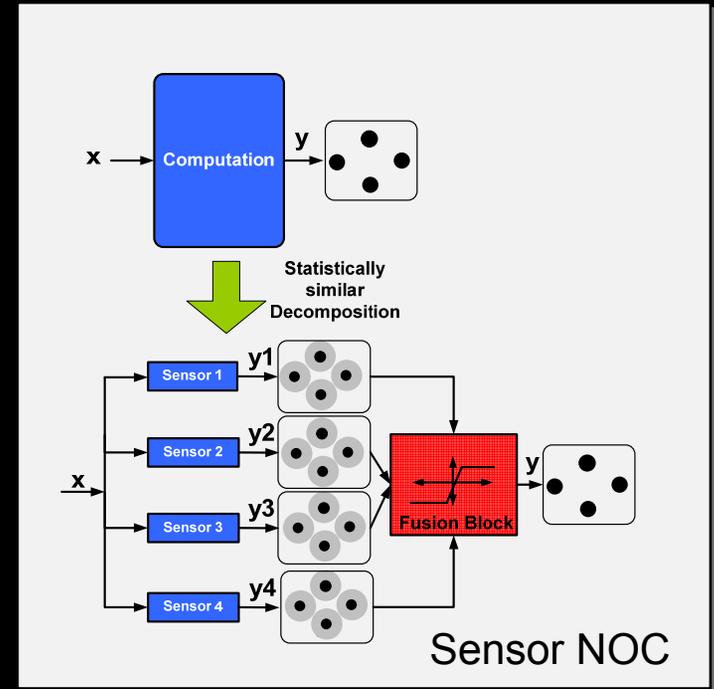
Carry Save Adder

A simple study:
2 different adders with
voltage over-scaling

- ◆ Architectural diversity makes error-error cross-correlation close to zero over range of frequencies and voltages
- ◆ Output SNR: CSA=2.9dB, DRCA=3dB, SNOC=5.5dB (2.5 dB improvement)

Example: PN code acquisition for CDMA

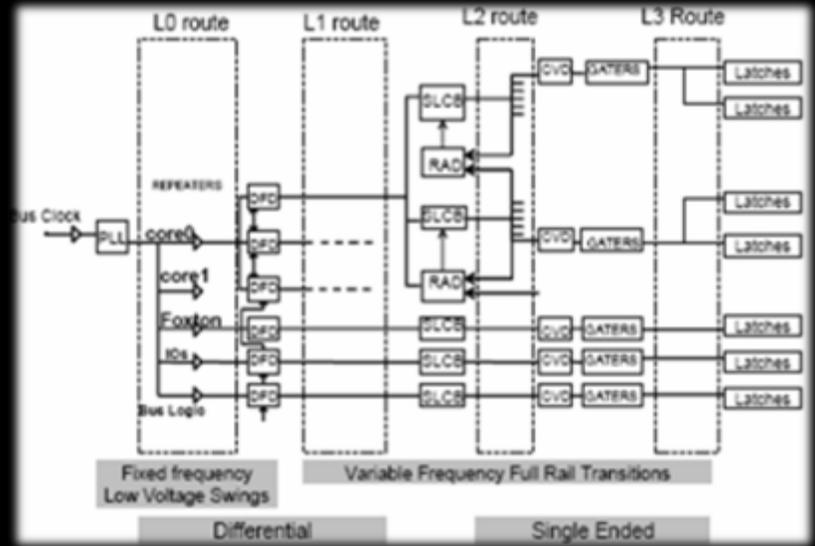
- *Statistically similar decomposition of function* for distributed sensor-based computation.
- Robust statistics framework for design of fusion block.
- Power savings of up to 40% for 8 sensors in PN-code acquisition in CDMA systems
- New applications in filtering, ME, DCT, FFT and others



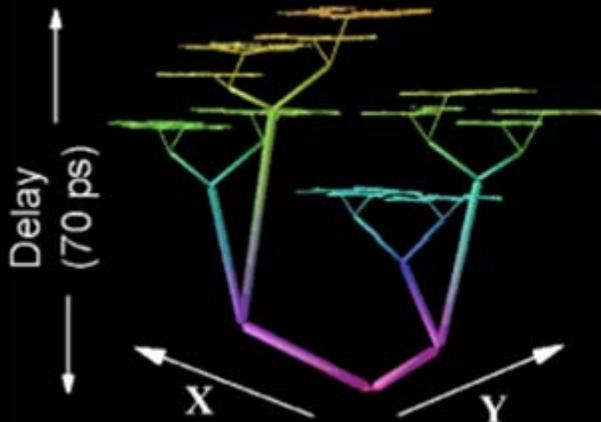
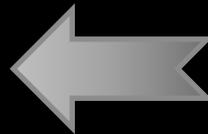
Example: State-of-the-art Synchronization



Precision Timing Element
(Crystal)



Intel Itanium Clock distribution
[ISSCC 05]



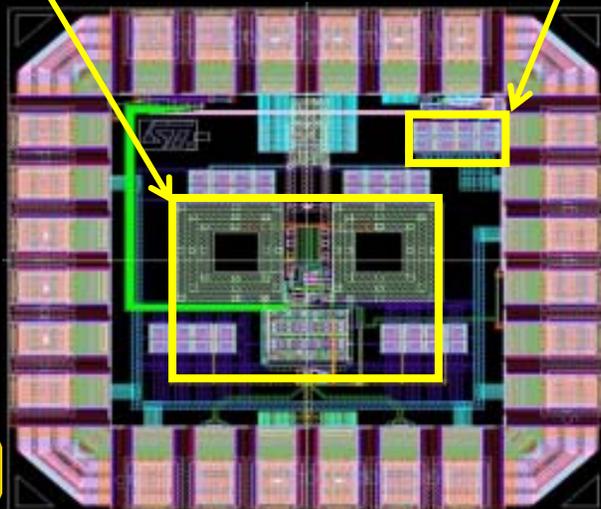
Clock phase and skew
[P. Restle, IBM]

Oscillators as Building Blocks

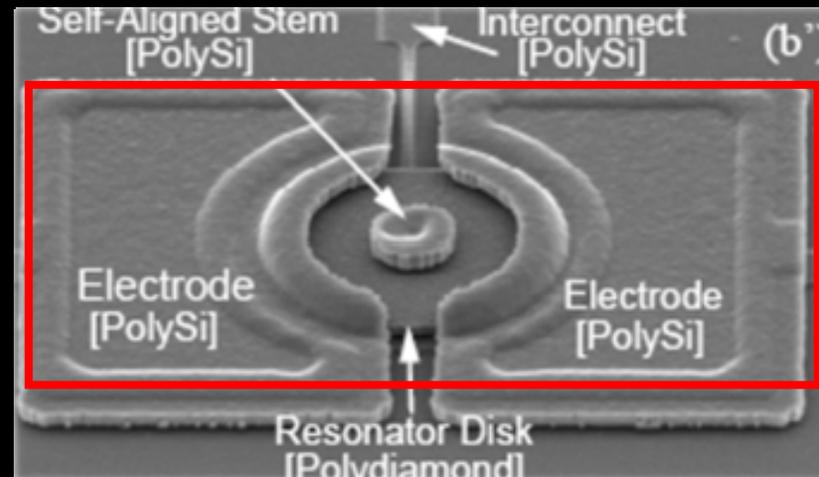
	Unit Area ($\mu\text{X}\mu$)		#/sq.mm	
LC			9	
		1 μW	750	7.5mW
	3x3	100 μW	90000	9W

LC Oscillator

Ring Oscillator



MEMS Disc Oscillator



Synchronization Inspired by Biological Systems

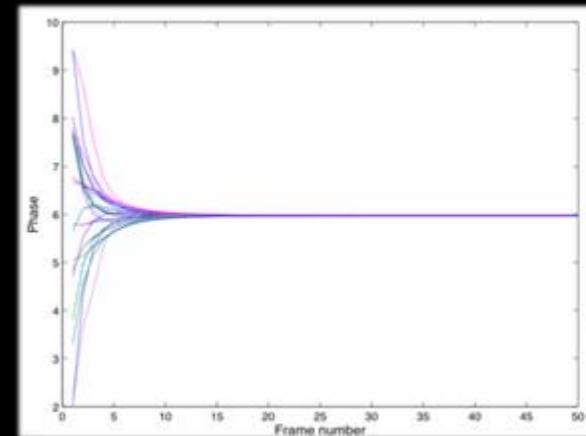
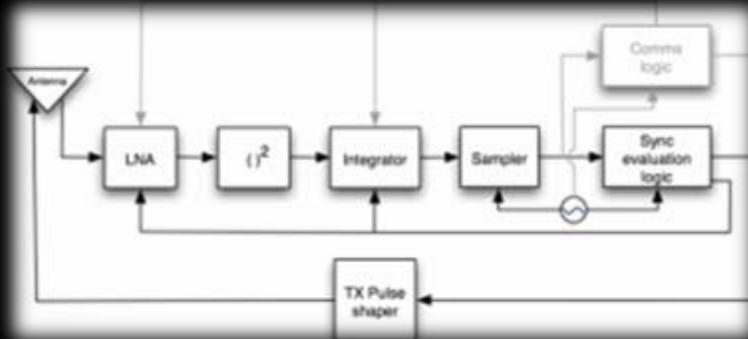
Distributed synchronization using only local communications and without precision timing elements



Energy distribution



time



Quick synchronization at low cost

Perspectives – Scaling The Wall

There is plenty of room at the bottom!

- Further scaling of energy/operation (or current per function) is essential for scaling to produce its maximum impact

- Current digital gates 5 orders of magnitude from minimum

- Two Major Take-Away's

- Always-optimal designs “park” themselves automatically in optimum energy point
- Aggressive deployments move beyond that point and use redundancy to recoup

It Takes A Systems Vision to Exploit the Offered Opportunities



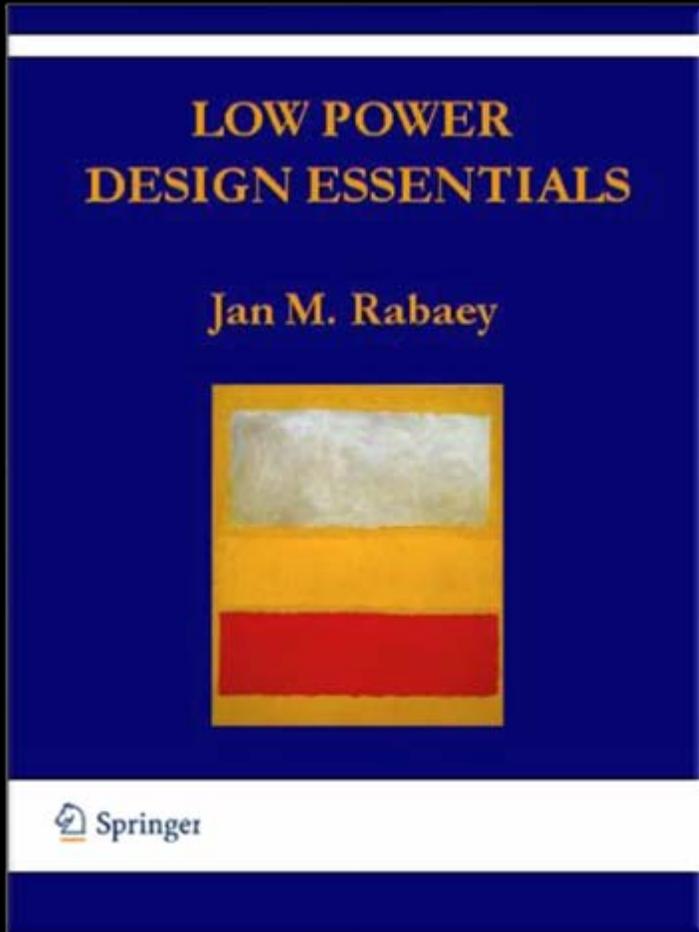
Thank you!

“Creativity is the ability to introduce order into the randomness of nature”

— Eric Hoffer

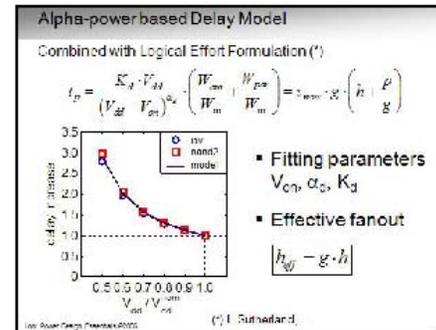
Acknowledgements: All of the GSRC and BWRC faculty and students, the funding by the FCRP and BWRC member companies and the US Government.

Expected by ISSCC 2008



TARGETING
EDUCATION AND
PROFESSIONAL
TRAINING

Innovative format

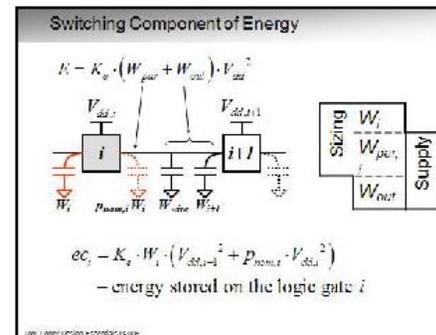


04.10

The delay of a logic gate is expressed using a simple linear delay model, based on the alpha-power law for the drain current. This is a curve-fitting expression and the parameters V_{th} and α_d are intrinsically related yet not equal to the transistor threshold and the velocity saturation index. K_d is another fitting parameter, and W 's correspond to various gate capacitances, with indices meaning output, parasitic and input. The model fits SPICE

simulated data quite nicely, across a range of supply voltages, normalized to the nominal supply voltage set by the technology, which is 1.2V in our case, for a 0.13 μ m CMOS technology.

Using the logical effort notation, the delay formula can be expressed simply as a product of the process-dependent time constant τ_{nom} and unit-less delay, where g is the logical effort that quantifies the relative ability of a gate to deliver current, h is the ratio of the total output to input capacitance, and ρ represents the delay component due to the self-loading of the gate. The product of the logical effort and the electrical effort is the effective fanout.



04.11

The energy of a logic gate is modeled by its switching component. In this models, $K_e W_{out}$ is the total load at the output, including wire and gate loads, and $K_e W_{par}$ is the self-loading of the gate. The total energy stored on these three capacitances is the energy taken out of the supply voltage in stage i .

On the other hand, if you change the size of the gate in stage i , it affects only the