Sub-45nm Circuit Technologies for High-performance Energy-efficient Microprocessors

Sanu Mathew and Ram Krishnamurthy
Circuit Research Lab, Intel Corporation
sanu.k.mathew@intel.com

Contributors: Mark Anders, Steven Hsu, Himanshu Kaul, Amit Agarwal, Yatin Hoskote, Nitin Borkar, Sapumal Wijeratne, Nanda Siddaiah, Bart Zeydel, Vojin Oklobdzija, David Harris, Wajdi Feghali, Kirk Yap

Sub-45nm Power-Performance Challenge

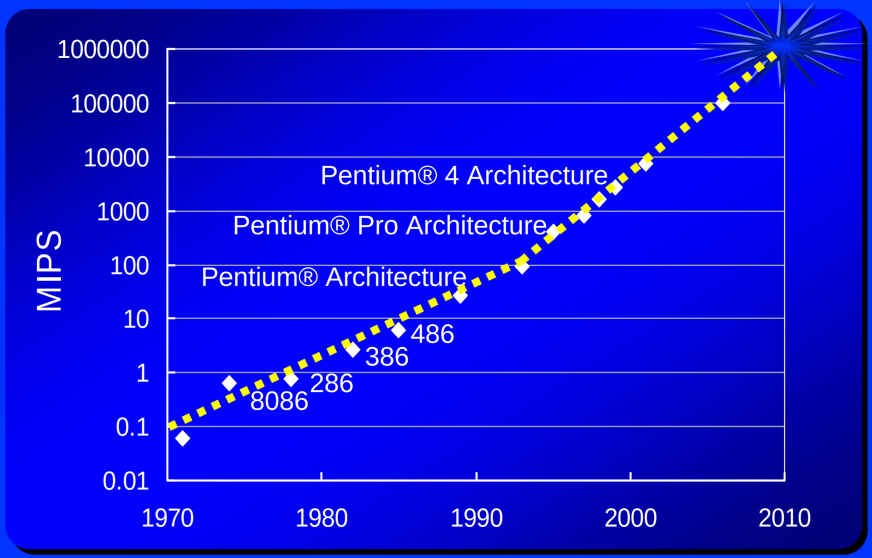

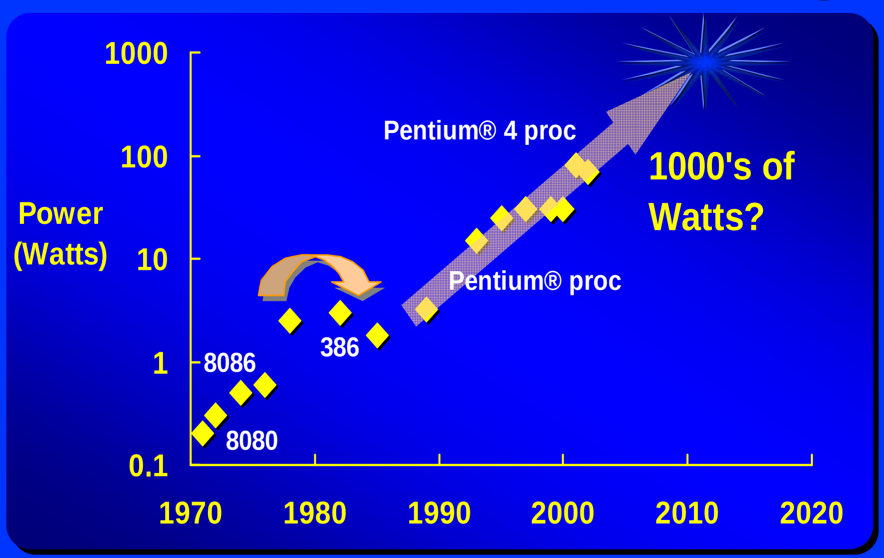

Strong demand for >TIPS performance in 2010+

Power will be the limiter to reach that

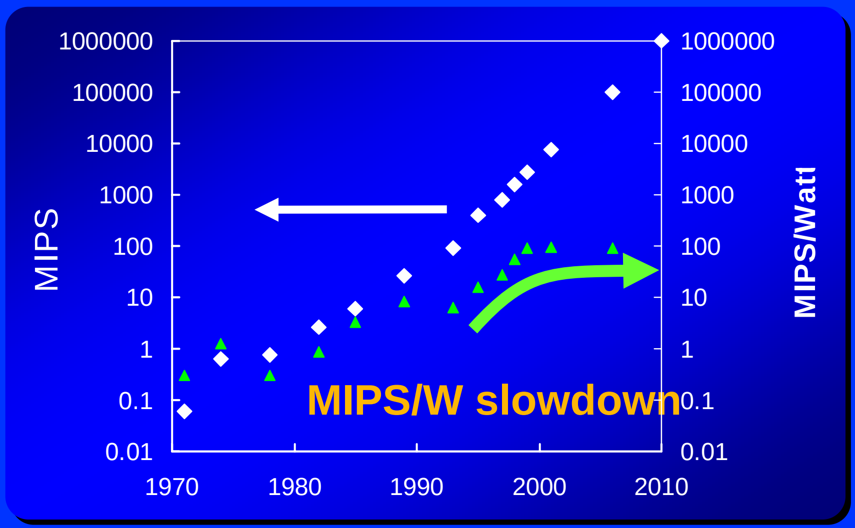

>2% power increase for every 1% performance ⇒ poor MIPS/Watt

Performance Through Parallelism

Platform 2015 Vision

“Over time, important functions once relegated to software and specialized chips are typically absorbed into the microprocessor itself. By moving functions on chip, such capabilities benefit from more efficient execution, superior economies of scale, and drastically reduced power consumption. <u>Special-purpose hardware is an important ingredient of Intel’s future processor and platform architectures</u>”.

Justin Rattner, “Platforms 2015”, IDF Keynote, March 3rd, 2005

GOPS/Watt Distinction: General-purpose vs. Dedicated

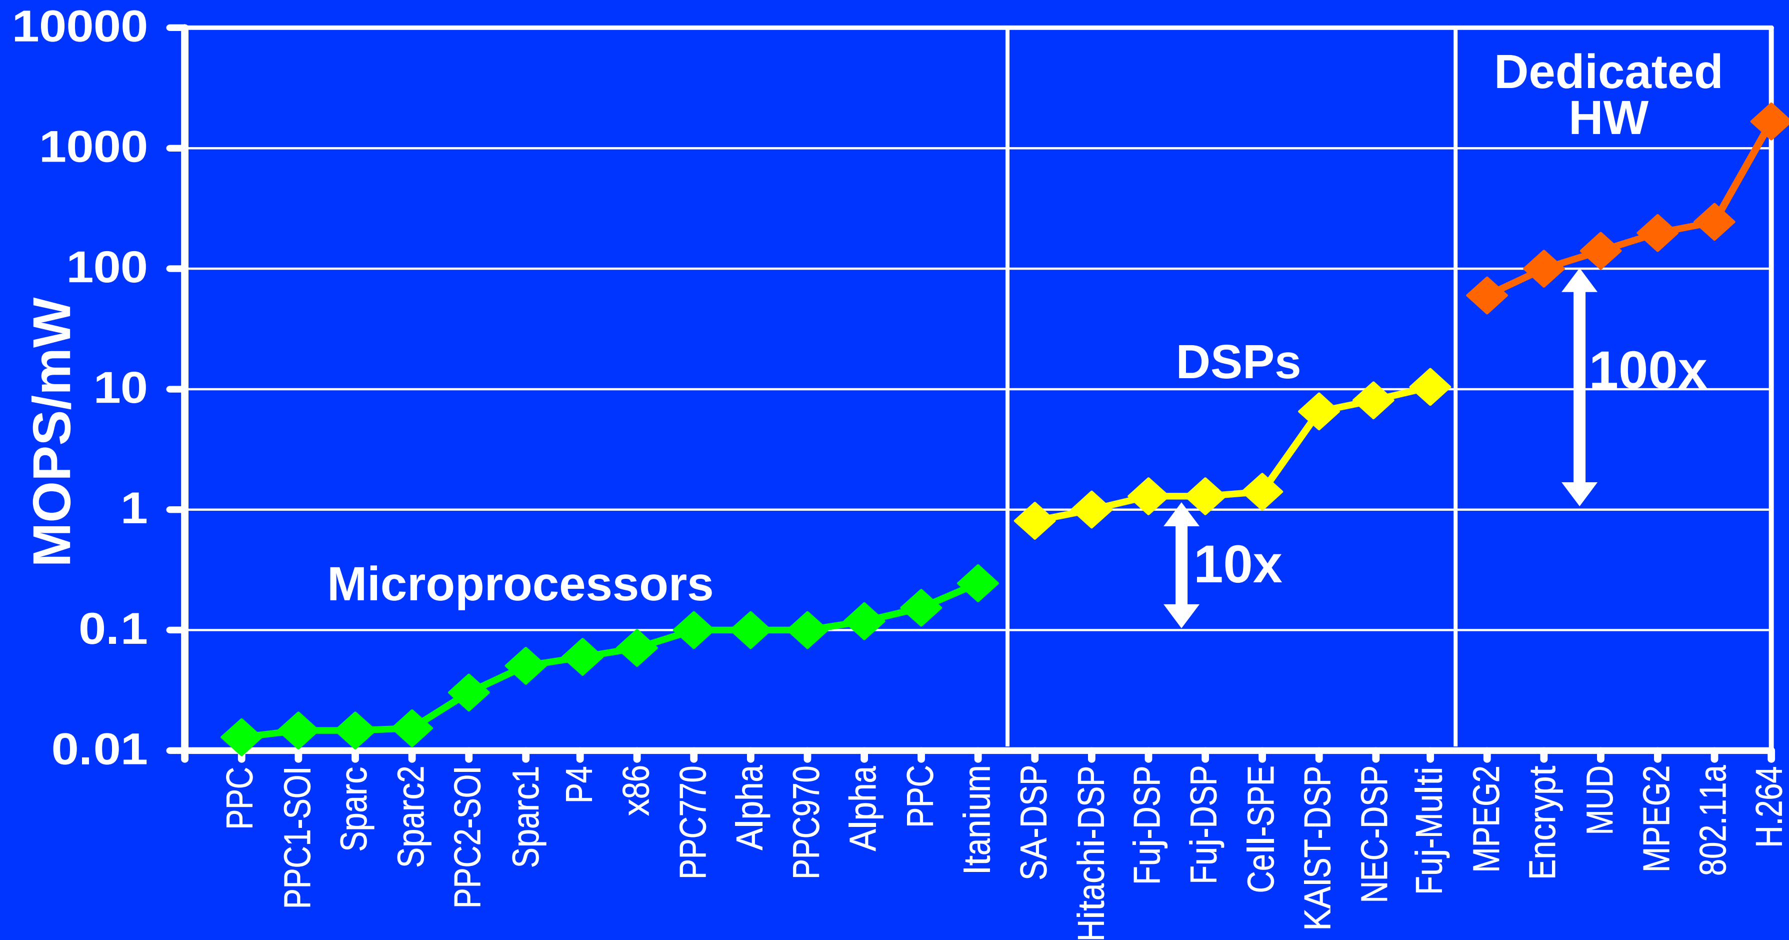

- Dedicated hardware: 100x higher energy-efficiency than GP
- DSP apps: Amenable to parallelism and pipelining

Efficient power-performance optimization

Special Purpose HW in Multi-core Processors

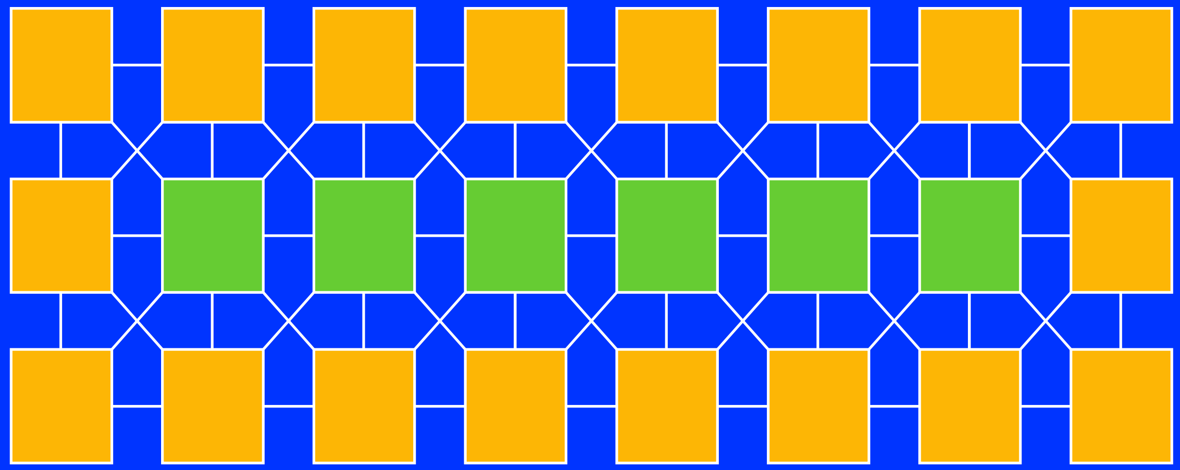

- **Low-power General-purpose core**
- **SP HW Accelerators**

1000GOPS+ SP HW accelerator cores integrated into a multi-core processor

Challenges:

- **Fixed function vs. limited programmable**
- **Multi/dual-supply voltage level converters**
- **Ultra low voltage operation, variation tolerance**

The Leap to Parallelism: Driving Energy-Efficient Performance

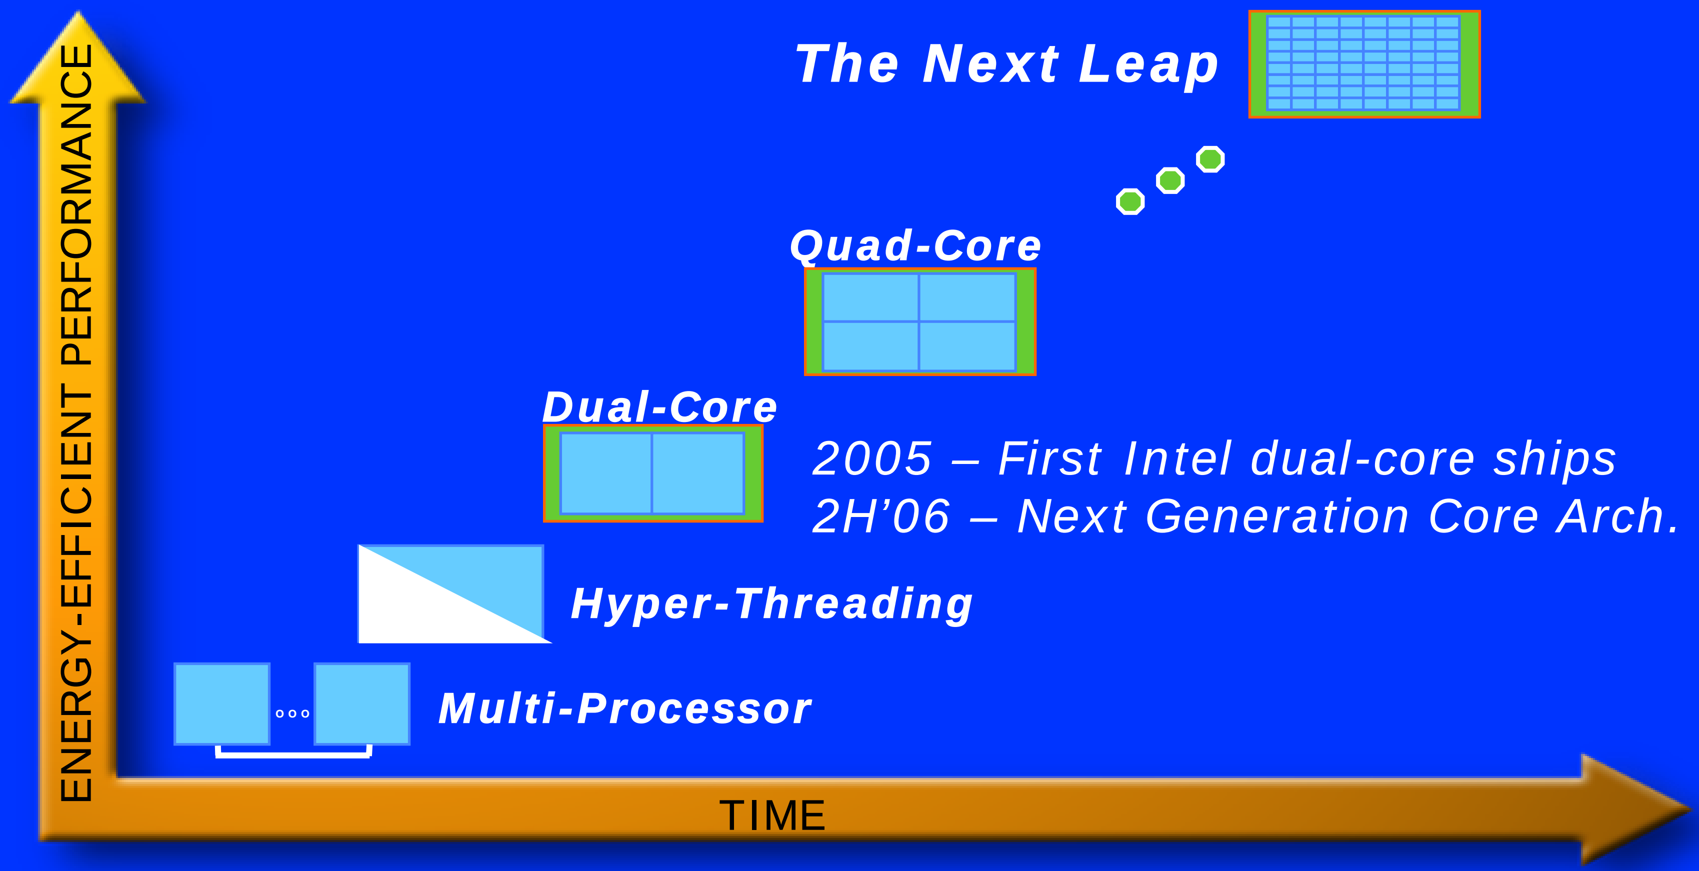

Special-purpose Hardware Accelerators: Next Leap in Performance beyond multi-core

Special-purpose HW for Viterbi Decoding

90nm 64-state Viterbi Accelerator

- 64-state radix-2 design: 40mW at 500Mbps in 90nm CMOS
- New leakage-tolerant trace-back register file and bit-serial ACS circuits
- **10X faster than current WLAN designs**

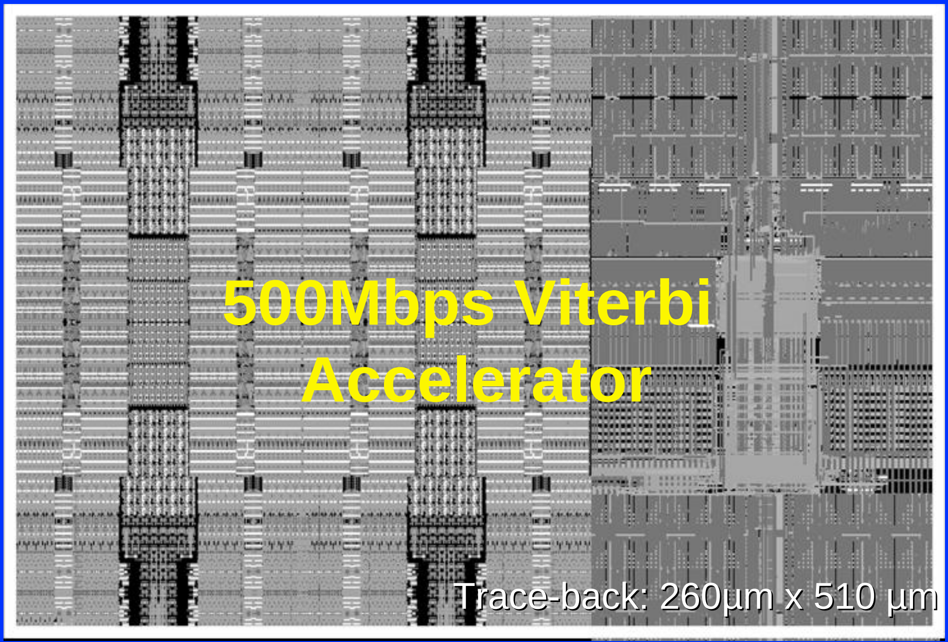

ACS: 230µm x 210 µm

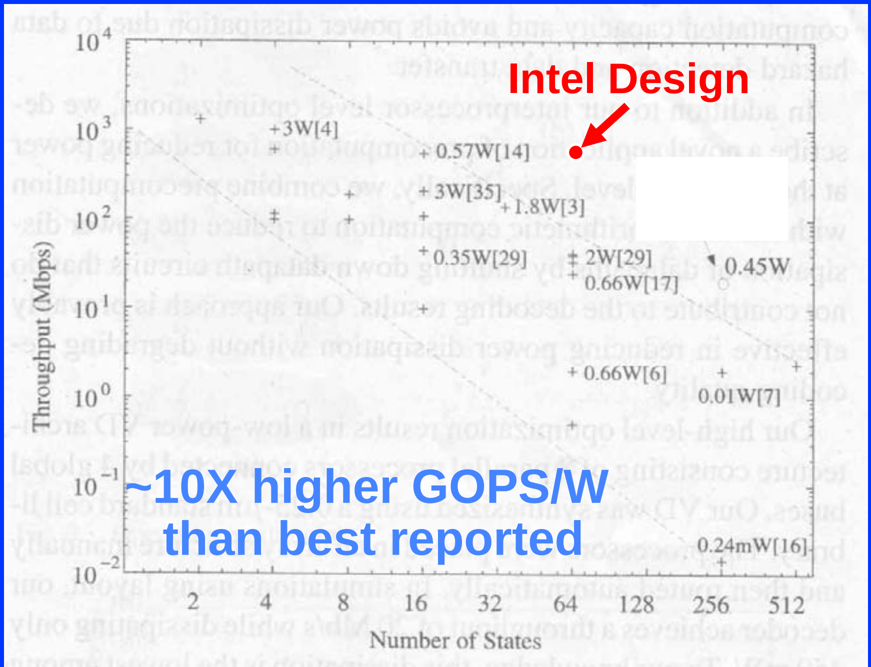

M. Anders et al, 2004 VLSI Circuits Symp.

Performance and power critical workload in wireless base-band, DVD codec, HDD signaling etc.

Special Purpose HW for DSP Filtering

90nm 110GOPS/Watt Filter Accelerator

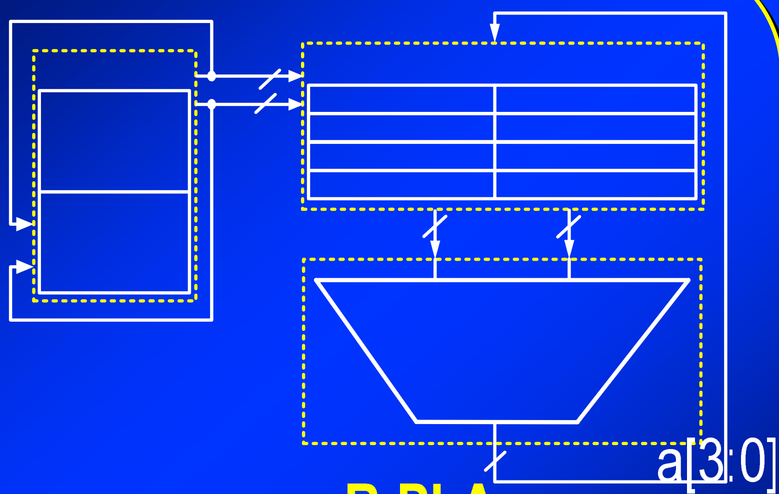

16-bit integer DSP multiplier with reconfigurable PLA control engine in 90nm CMOS

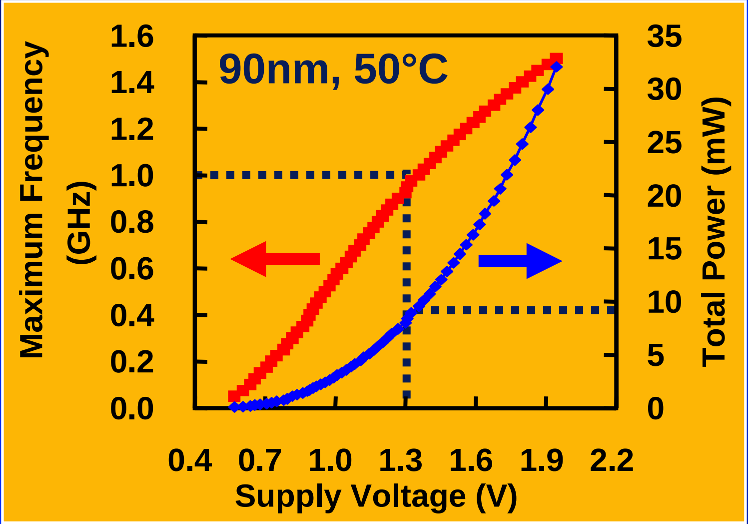

1GHz single-cycle throughput at 9mW

Highest reported performance per watt

>3X benefit over previously published design

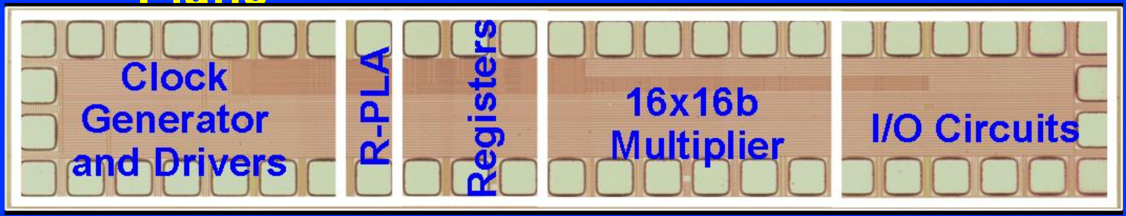

S. Hsu et al, ISSCC'05

Performance and power critical workload in video, graphics, SIMD, base-band (FIR, FFT, DCT)

Filter Accelerator: Comparisons

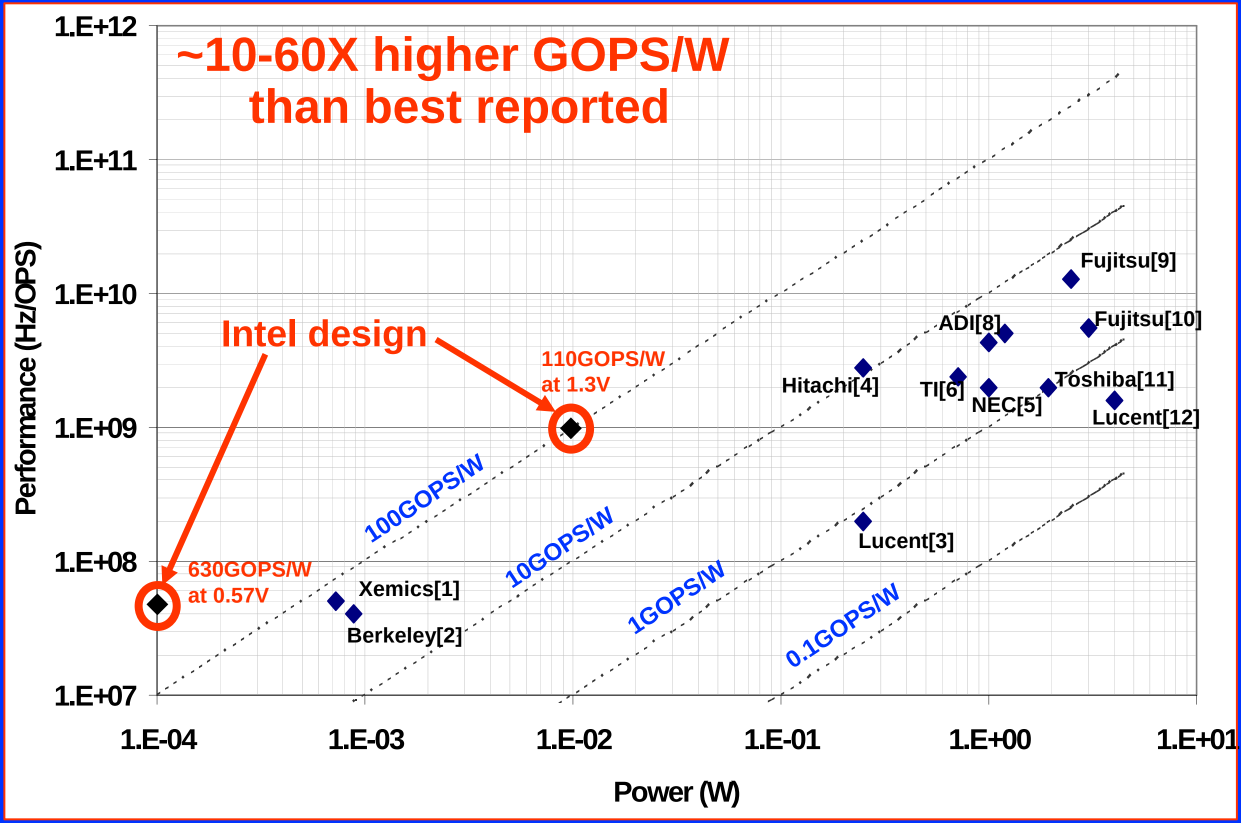

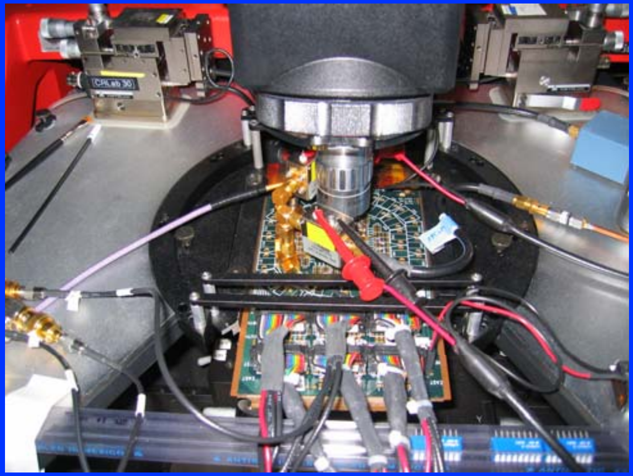

Measurement test setup

9pJ/OP or 110GOPS/Watt at nominal 1.3V

1.6pJ/OP or 630GOPS/Watt at 0.57V

Highest Filter MACC power-performance in industry

Crypto Accelerator

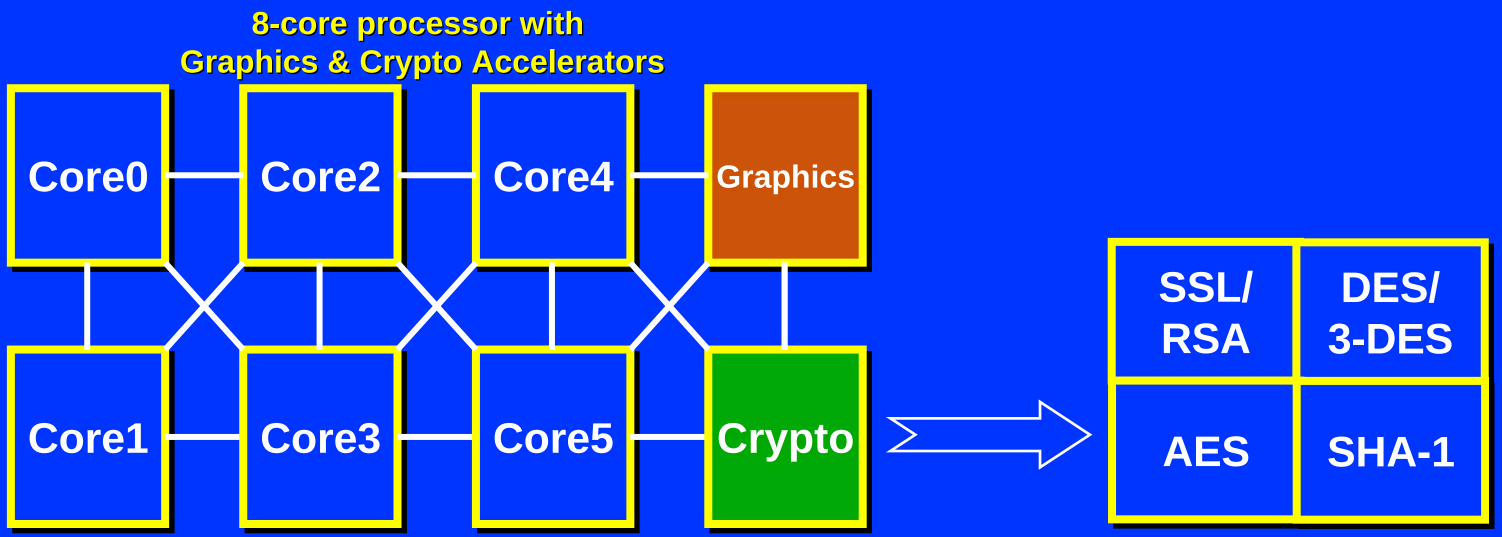

- Encryption apps are computationally intensive due to:
  - Wide operand bit-widths (eg: 2048-bit modular exponentiations)
  - Iterative, parallelizable algorithms (eg: 48-round bit permutations)
- Energy-inefficient on GP execution cores
- Hardware accelerator offers 5x higher performance/watt
- Encryption workloads offloaded to a Crypto-accelerator

Crypto Accelerator: SSL

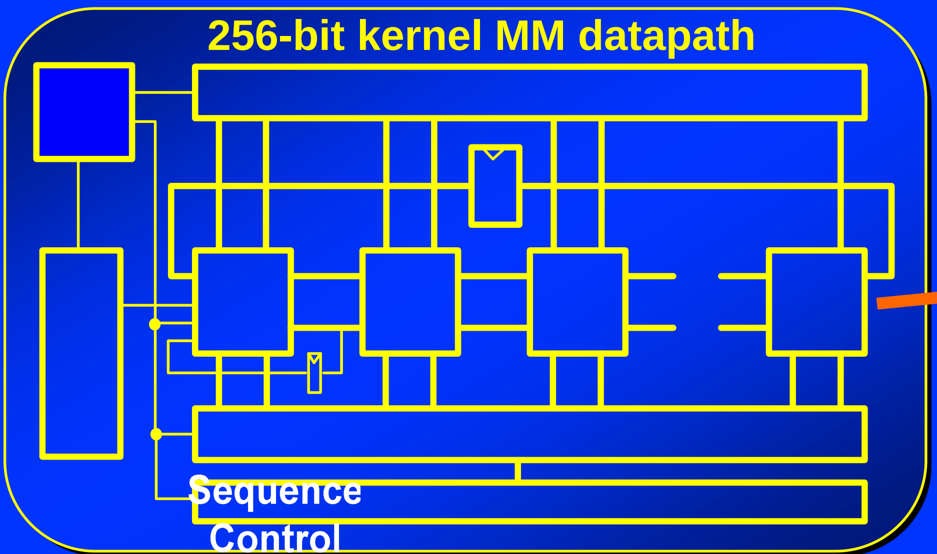

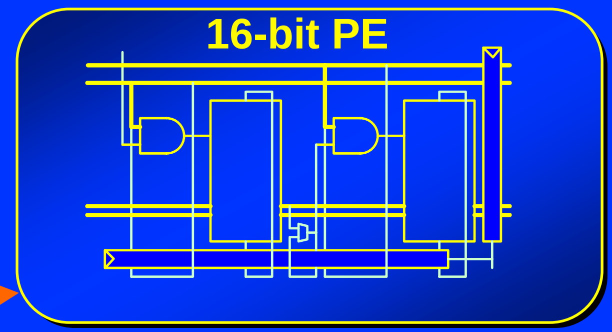

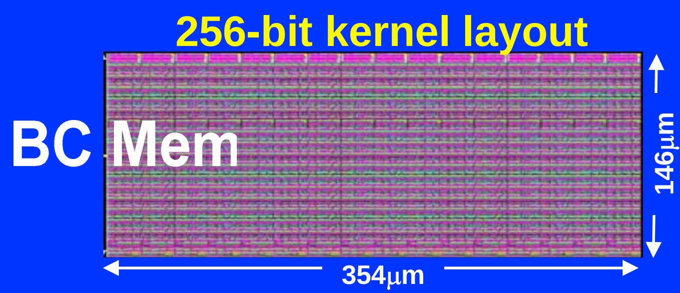

D. Harris, S. Mathew et al, 2005 Arith-17 Symp.

- **A×B *mod* C is a key operation in RSA cryptography**
- **Scalable design ⇒reconfigurable for 256/1024bit op.**
- **Reconfigurable for RSA, Diffie-Helman, DSA, ECC**
- **15K 256-bit exponentiations/s and 7.3MMults/s**
- **44% speedup over prior-art**

PE$_{16}$

Z Mem

Special Purpose HW for TCP-processing

TCP Offload Engine

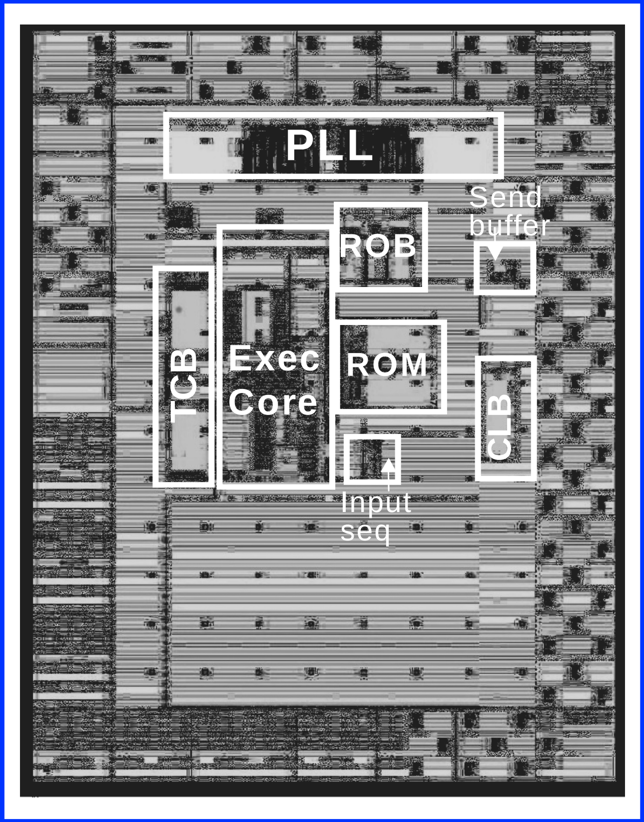

2.23 mm X 3.54 mm, 260K transistors
(Y. Hoskote, et. al. ISSCC '03)

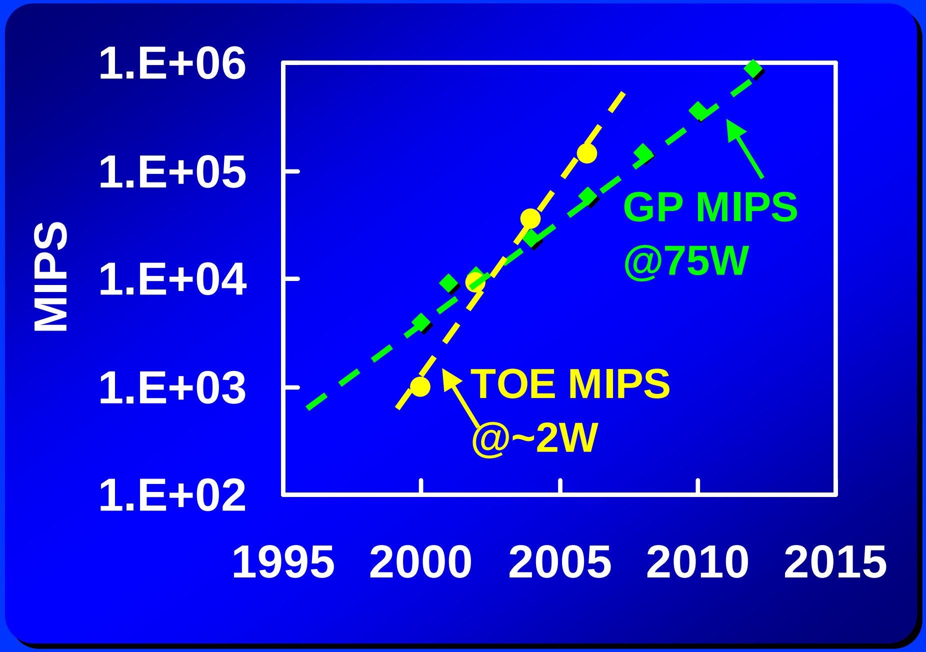

Opportunities for acceleration:
Network processing engines
MPEG Encode/Decode engines
Speech engines

Special purpose HW—Best MIPS/Watt

GP Processor Core Challenges

- Energy-efficient ALU
  - Sparse-tree adder core
- Low-power dual-V_{cc} operation
  - Split-output level converter
- High-performance interconnects
  - Full-swing transition-encoded domino busses
- Leakage/Variation tolerant memories
  - Full-swing register file with compensation keeper

GP Processor Core Challenges

- Energy-efficient ALU
  - Sparse-tree adder core
- Low-power dual-V_{cc} operation
  - Split-output level converter
- High-performance interconnects
  - Full-swing transition-encoded domino busses
- Leakage/Variation tolerant memories
  - Full-swing register file with compensation keeper

Execution Core Power Density

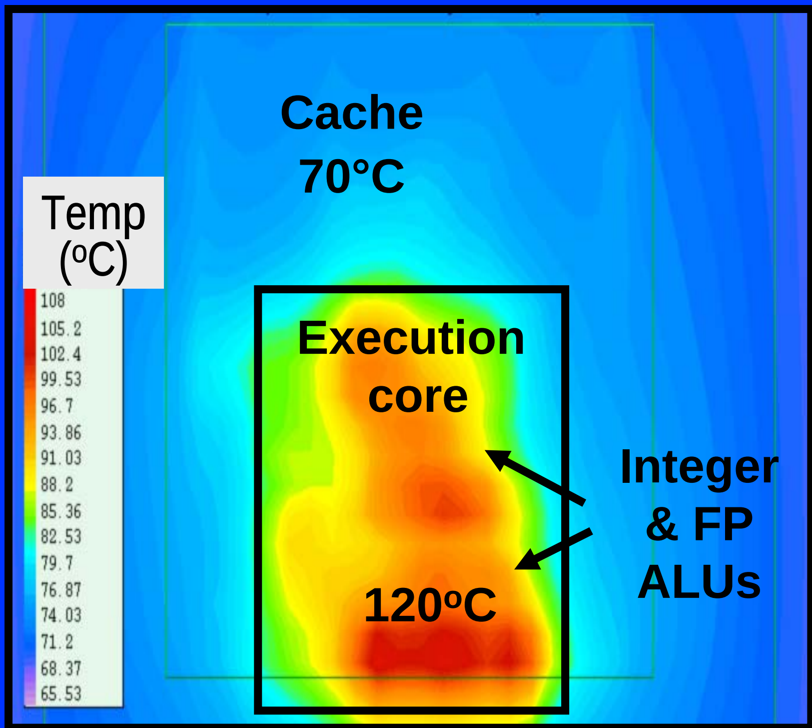

Itanium2® thermal map

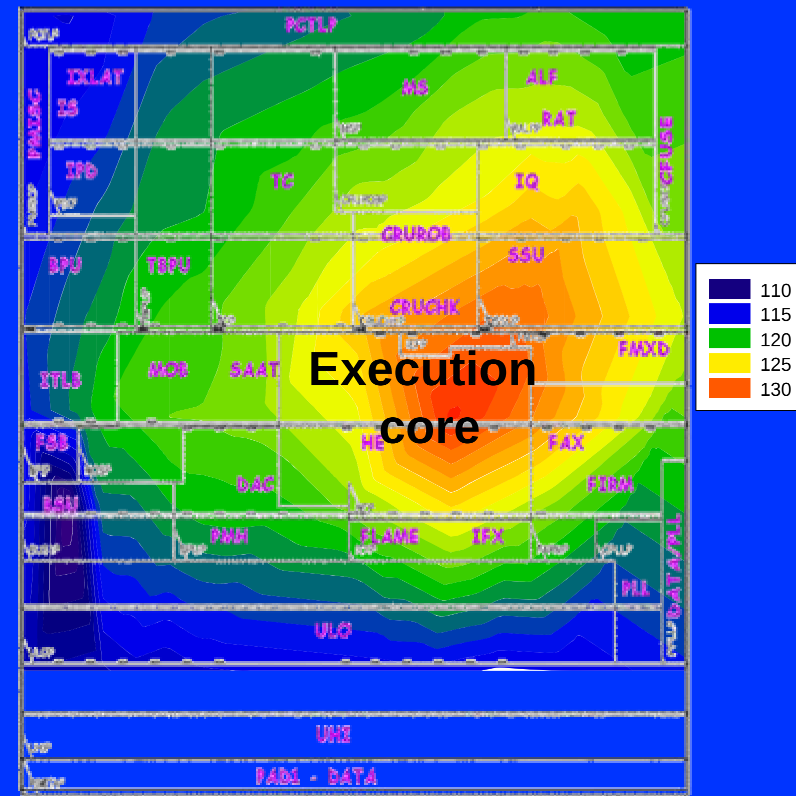

Pentium4® thermal map

- Integer & FPU ALUs/AGUs: performance & power limiters
- High activity ⇒ thermal hotspots and peak-current limiters

Requires high-performance, low-power execution core circuits

Sparse-Tree Adders

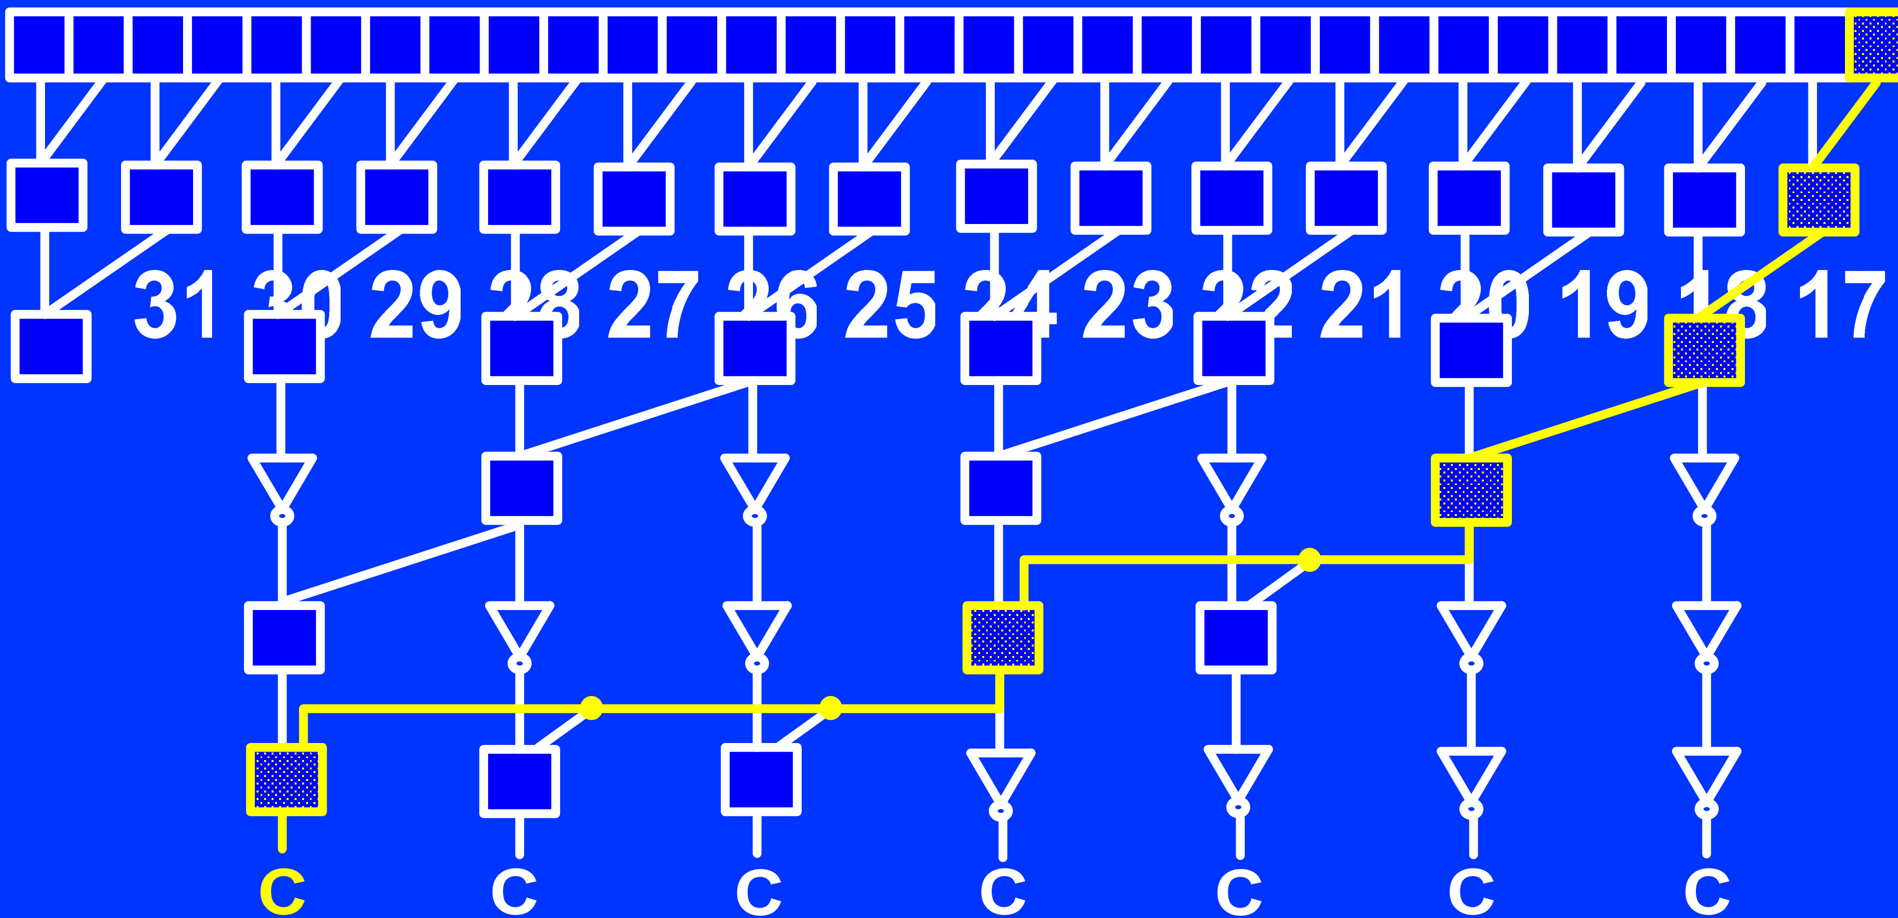

- Generate only every 4th carry in parallel, i.e., C_3, C_7 etc.
- Non-critical 4-bit sum generator side path to produce sum
- 73% fewer carry-merge gates ⇒ fast and energy-efficient

Non-critical Conditional Sum Generator

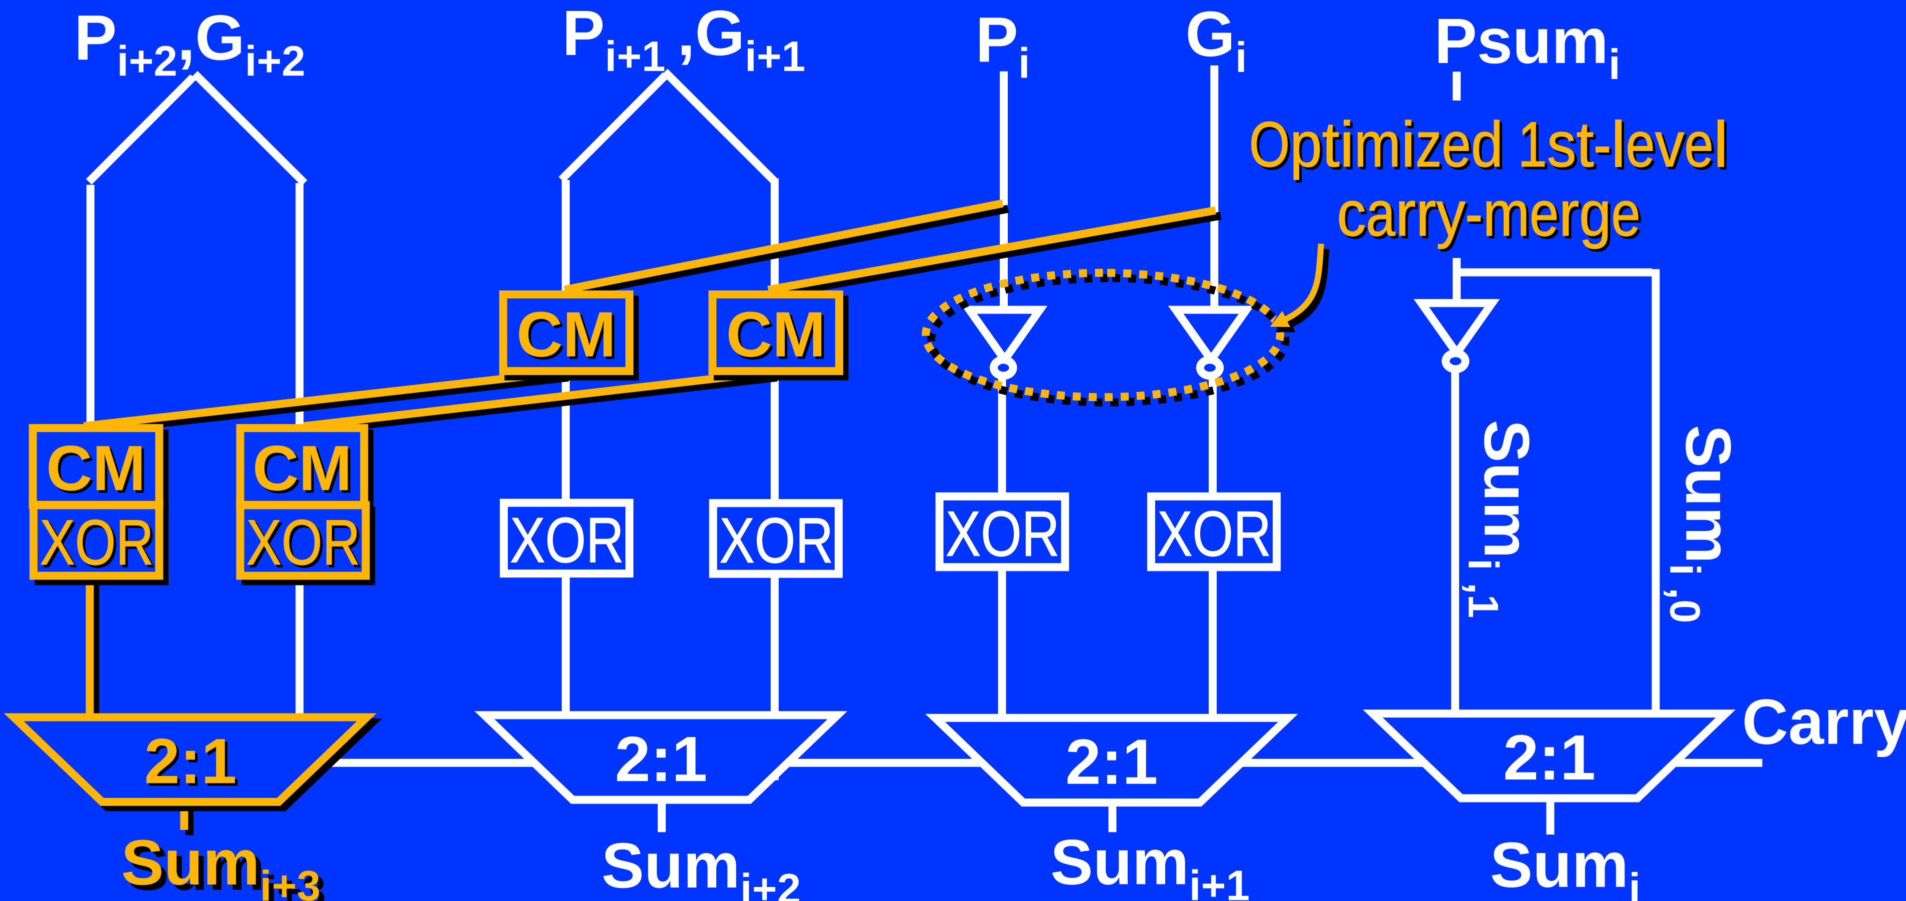

- Non-critical path: 4-bit ripple carry chain
  - Reduced area, energy consumption, and active leakage
- Generates conditional sums for each cluster of 4-bits
- Sparse-tree carry selects appropriate sum

90nm 7GHz 64-bit Integer ALU (ISSCC'04)

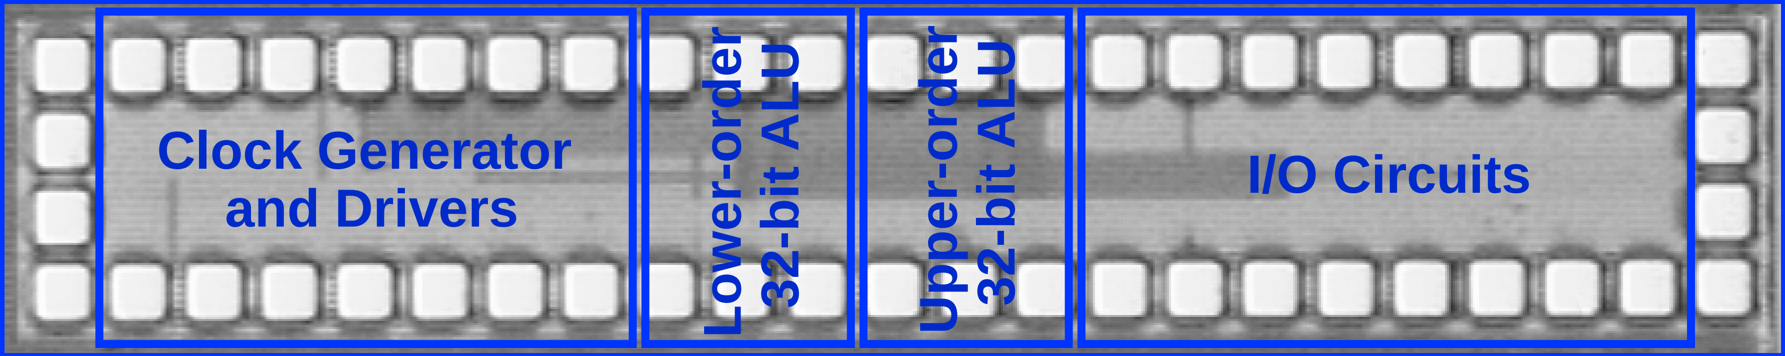

S. Mathew et al,
ISSCC 2004 & JSSC 01/05

Process	90nm Dual-Vt CMOS, 7 Metal
Die area	0.474mm^2
64-bit ALU layout area	0.073mm^2
Total transistor count	6100
64-bit ALU average switching power (α=0.3)	89mW at 4GHz, 1.3V, 25ºC
64-bit ALU active leakage power	9.6mW at 1.3V, 25ºC
64-bit ALU maximum frequency	7GHz at 2.1V, 25C
32-bit ALU average switching power (α=0.3)	71mW at 7GHz, 1.3V, 25ºC
32-bit ALU active leakage power	4.4mW at 1.3V, 25ºC

64-bit ALU die microphotograph and measured performance summary

- 7GHz single-cycle 64-bit integer ALU (measured in 90nm CMOS)
- Simultaneous 9GHz single-cycle 32-bit integer ALU mode

Fastest reported single-cycle 64-bit integer ALU performance

GP Processor Core Challenges

- Energy-efficient ALU
  - Sparse-tree adder core
- Low-power dual-V_{cc} operation
  - Split-output level converter
- High-performance interconnects
  - Full-swing transition-encoded domino busses
- Leakage/Variation tolerant memories
  - Full-swing register file with compensation keeper

Split-output Level Converter

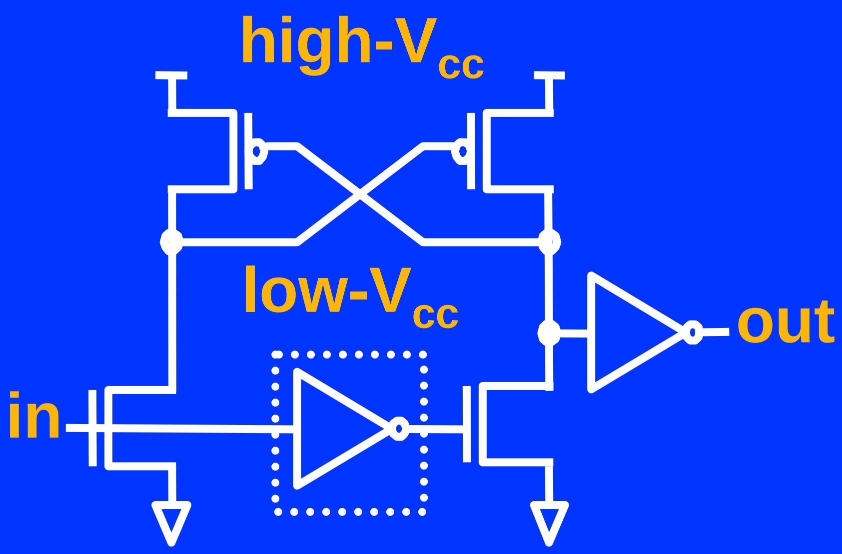

Conventional CVSL LCB

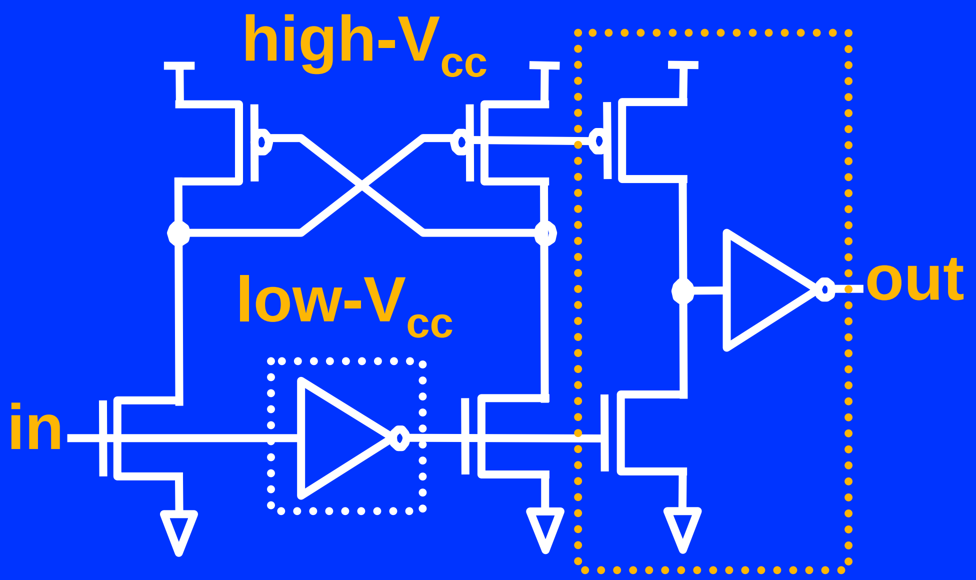

Split-output LCB

- Contention in CVSL LCB degrades delay
- Split-output LCB decouples CVSL stage from output driver stage
  - Fast level conversion due to low contention
  - Reduced fanin load on clock grid

LCB Energy-Delay Comparisons

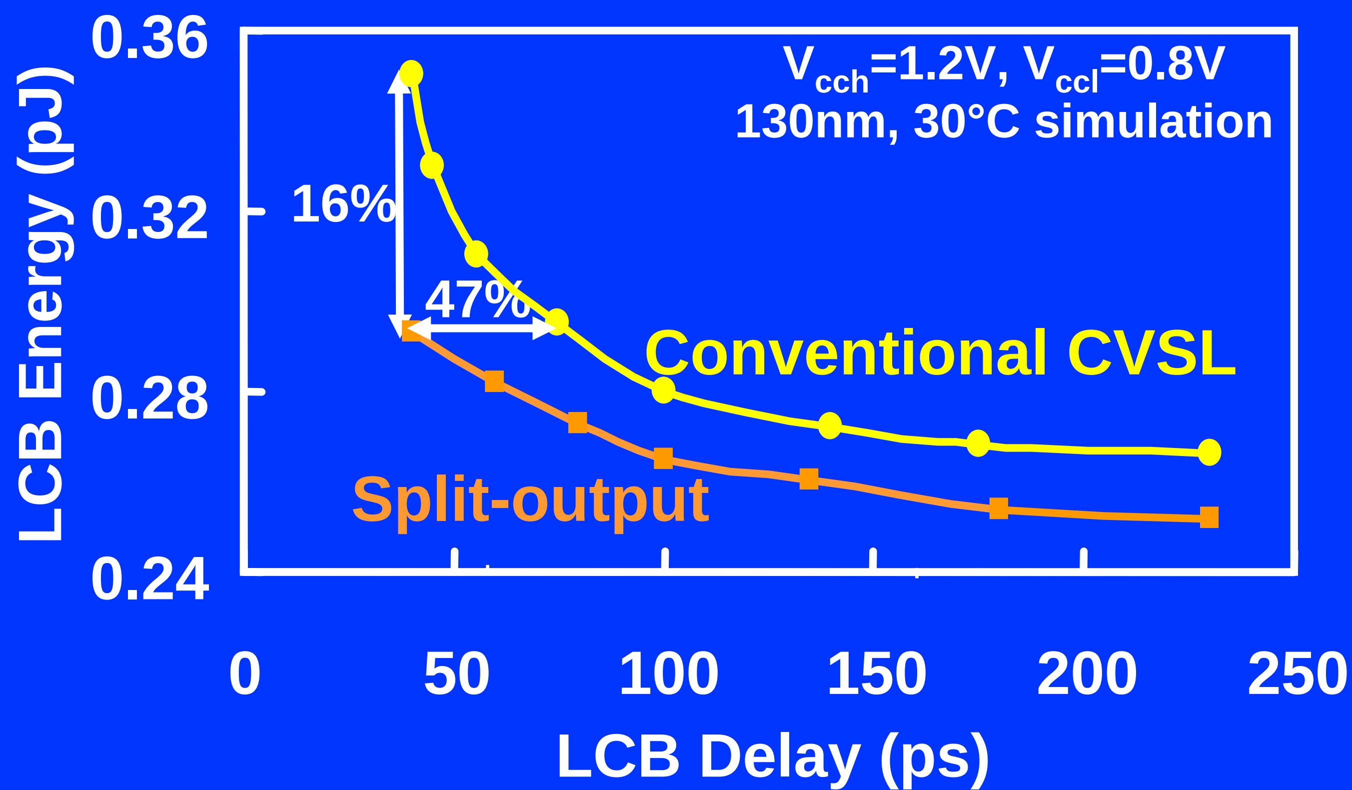

R. Krishnamurthy et al, 2002 Symp. VLSI Circuits

LCB Scheme	Fanin cap (fF)	Total area (mm²)	CVSL-stage contention energy (pJ)
Conventional CVSL	8.2	15.5	0.085
This work	7.1 (-14%)	13.8 (-11%)	0.039 (-54%)

Effective low-energy alternative to CVSL LCB

GP Processor Core Challenges

- Energy-efficient ALU
  - Sparse-tree adder core
- Low-power dual-V_{cc} operation
  - Split-output level converter
- High-performance interconnects
  - Full-swing transition-encoded domino busses
- Leakage/Variation tolerant memories
  - Full-swing register file with compensation keeper

Transition-Encoded Bus

M. Anders et al, 2002 VLSI Circuits Symp.

- **Encoder circuit**
  - **XOR of previous and current input**
- **Decoder circuit**
  - **XOR of previous output and bus state**
- **Domino delay performance**
  - **Collinear cap reduction**
- **Static bus energy**
  - **Transition-dependent activity**

Transition-Encoded Bus: Results

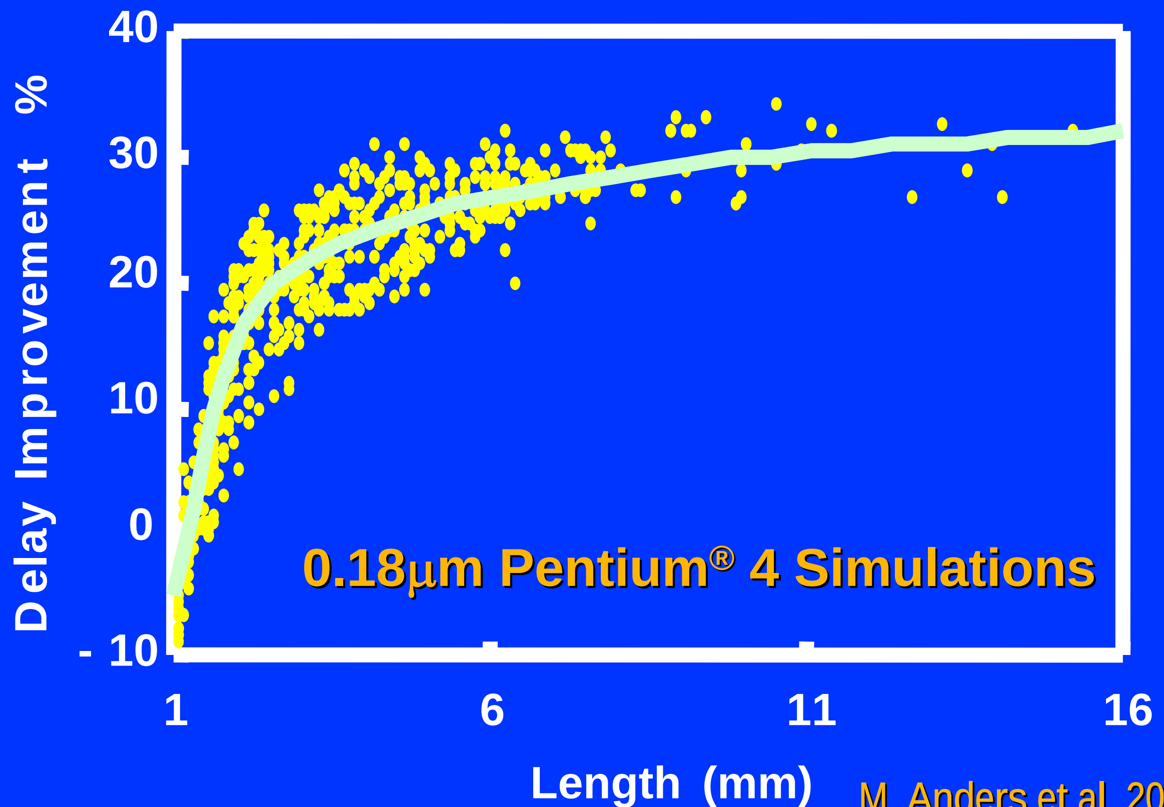

M. Anders et al, 2002 VLSI Circuits Symp.

- Transition only when current input != previous input
- Dynamic bus performance but energy profile of static bus
- Energy scales linearly with input switching activity
- 79% of full-chip buses: 10%-35% delay improvement

GP Processor Core Challenges

- Energy-efficient ALU
  - Sparse-tree adder core
- Low-power dual-V_{cc} clocking
  - Split-output level converter
- High-performance interconnects
  - Full-swing transition-encoded domino busses
- Leakage/Variation tolerant memories
  - Full-swing register file with compensation keeper

PVT Variation tolerant memories: Keeper Upsizing

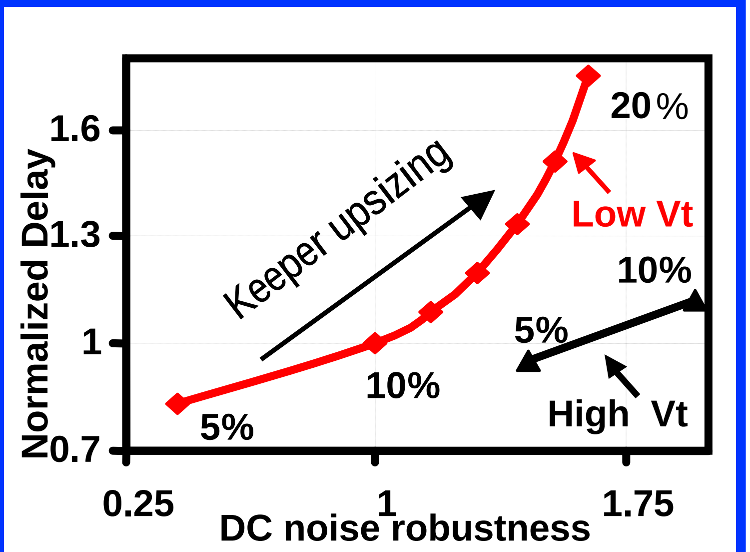

A. Alvandpour et al, 2001 Symp. VLSI Circuits

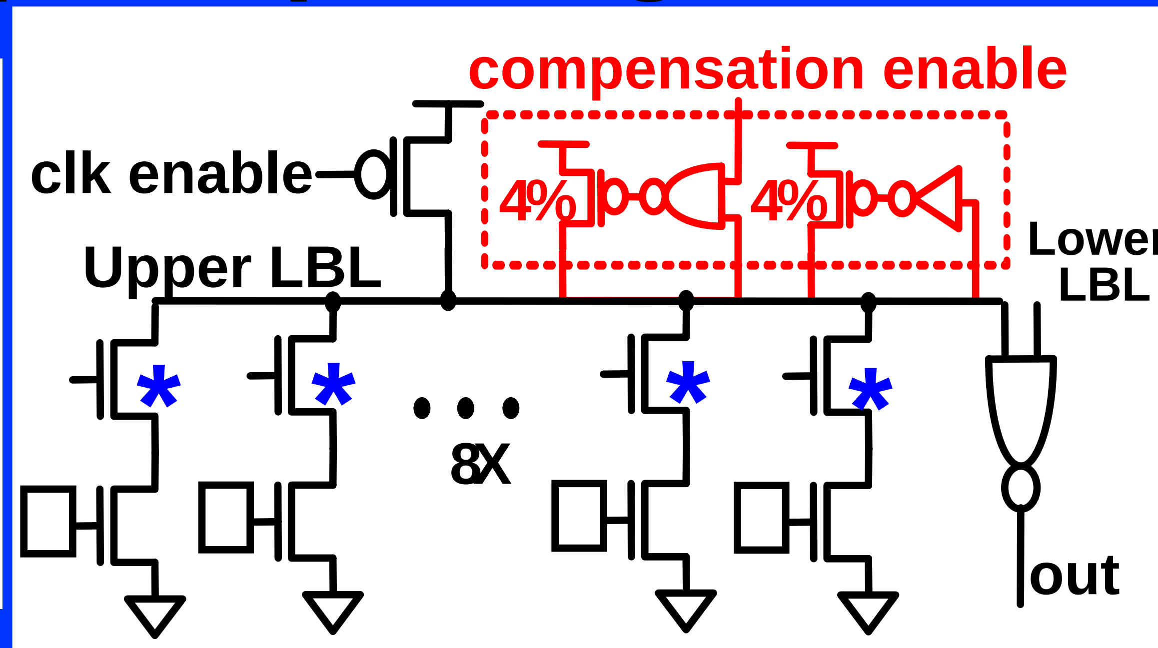

- PVT induced leakage variation
  - Traditional noise engineering: diminishing ROI
- Additional keeper enables customization
  - High leakage states: 8% keeper
  - Low leakage states: 4% keeper

PVT Compensation Motivation

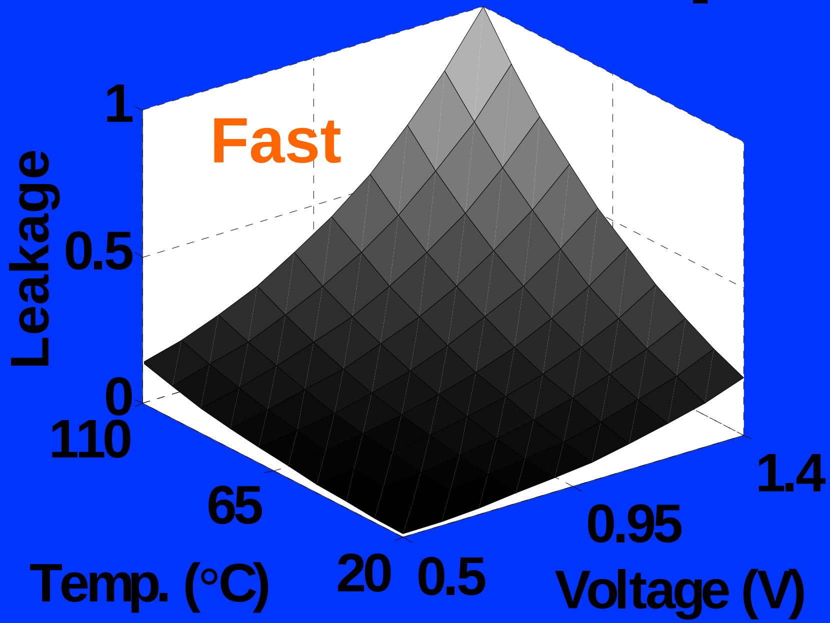

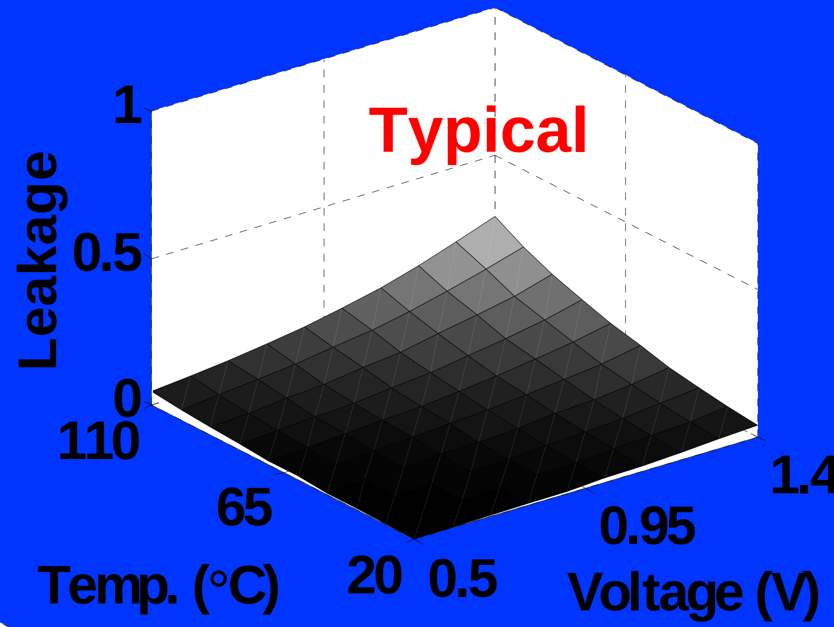

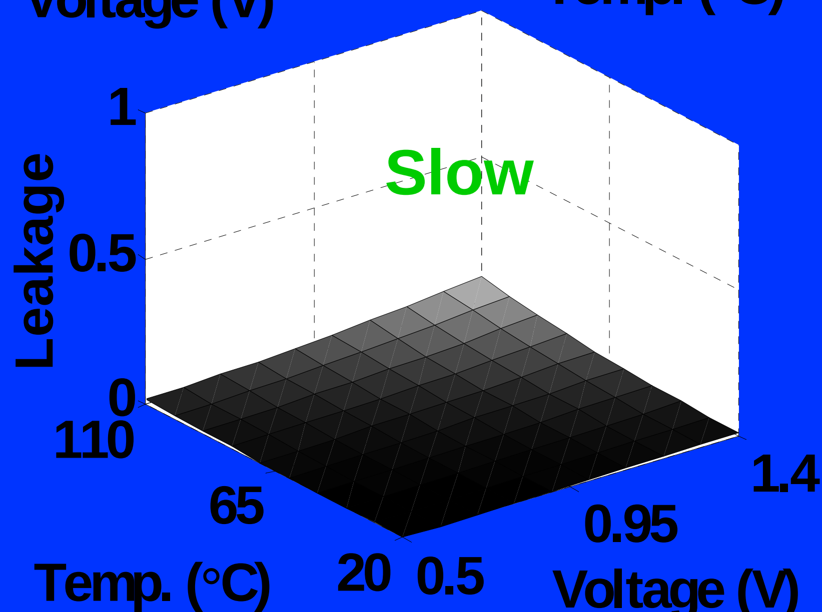

10^4X leakage variation across PVT

PVT Compensation Benefit

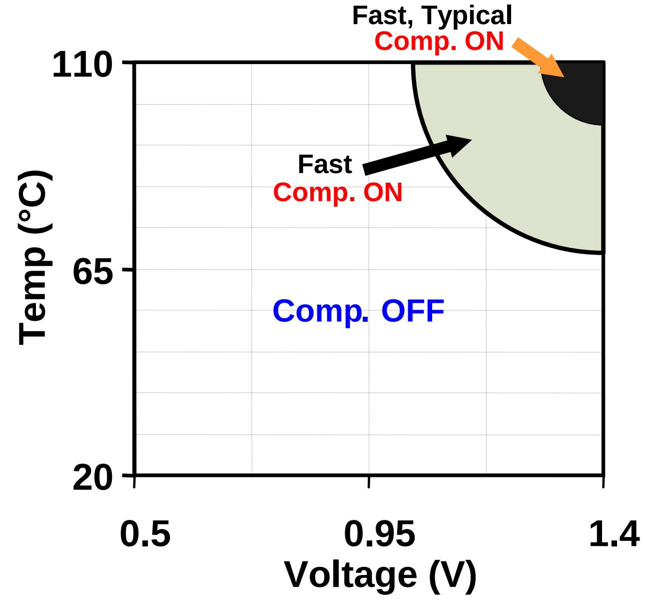

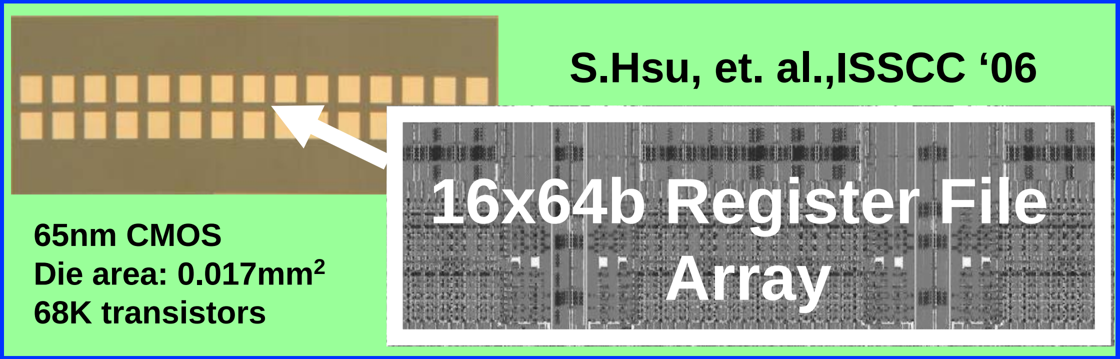

P1264 Si measurements, 50°C

	Frequency	Power	Leakage
Nominal (1.2V)	8.8GHz	198mW	25mW
Peak Performance (1.4V)	10.1GHz	273mW	57mW
Low-Voltage Mode (0.5V)	300MHz	1.3mW	405μW

- High leakage dies: 27% robustness ↑
- Low leakage dies: 10% delay ↓

Industry's highest performance/Watt ⇒ 96GOPs/Watt

Summary

- Performance through parallelism: Multi-core
- Special-purpose hardware accelerators provide higher MIPS/Watt vs. general-purpose cores
- Energy-efficient, leakage/variation tolerant circuits required for scalable GP performance
  - Sparse-tree adder
  - Split-output level converter
  - Transition-encoded interconnects
  - Leakage/variation tolerant register files