An IA-32 Processor with a Wide Voltage Operating Range in 32nm CMOS

Gregory Ruhl, Saurabh Dighe, Shailendra Jain, Surhud Khare, Satish Yada, Ambili V, Praveen Salihundam, Shiva Ramani, Sriram Muthukumar, Srinivasan M, Arun Kumar, Shasi Kumar, Rajaraman Ramanarayanan, Vasantha Erraguntla, Jason Howard, Sriram Vangal, Paolo Aseron, Howard Wilson, Nitin Borkar

Microprocessor & Programming Research, Intel Labs

Purpose

Claremont: Near Threshold Voltage and Wide Dynamic Range IA Core

Demonstrate energy benefits of Near Threshold Voltage (NTV) computing to IA

Extend dynamic range of operation from NTV to Vmax for energy efficient performance

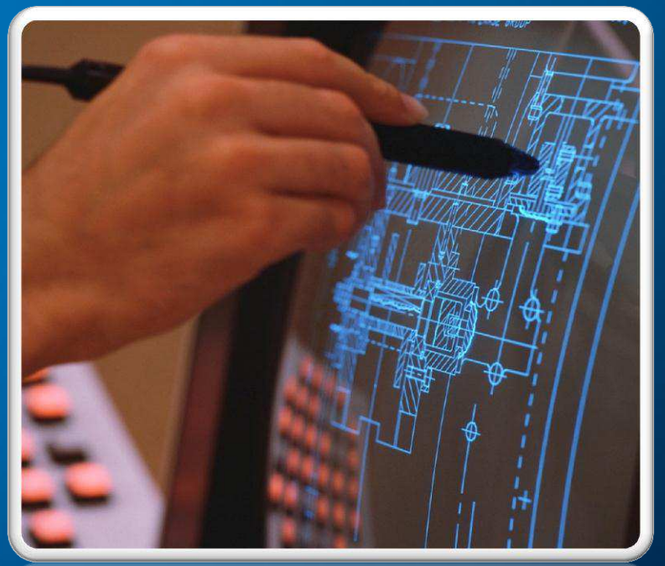

Advance low voltage, variation aware and multi-corner design methodologies

Agenda

- Design Challenges
- Claremont Prototype
- Design Strategies and Methodologies
  - NTV Design
  - Wide Dynamic Range Design
- Results and Summary

Design Challenges

- Reduced I_{on}/I_{off}, noise margins and variability results in circuit functional failures
- Power/performance profile becomes extremely sensitive to PVT variations
- Tools and methodologies are not mature for low voltage designs
- Wide dynamic range design convergence is complicated by
  - Disproportionate device vs. interconnect delay scaling
  - Multiple voltage domains

Agenda

- Design Challenges
- Claremont Prototype
- Design Strategies and Methodologies
  - NTV Design
  - Wide Dynamic Range Design
- Results and Summary

P54C IA-32 Core Background

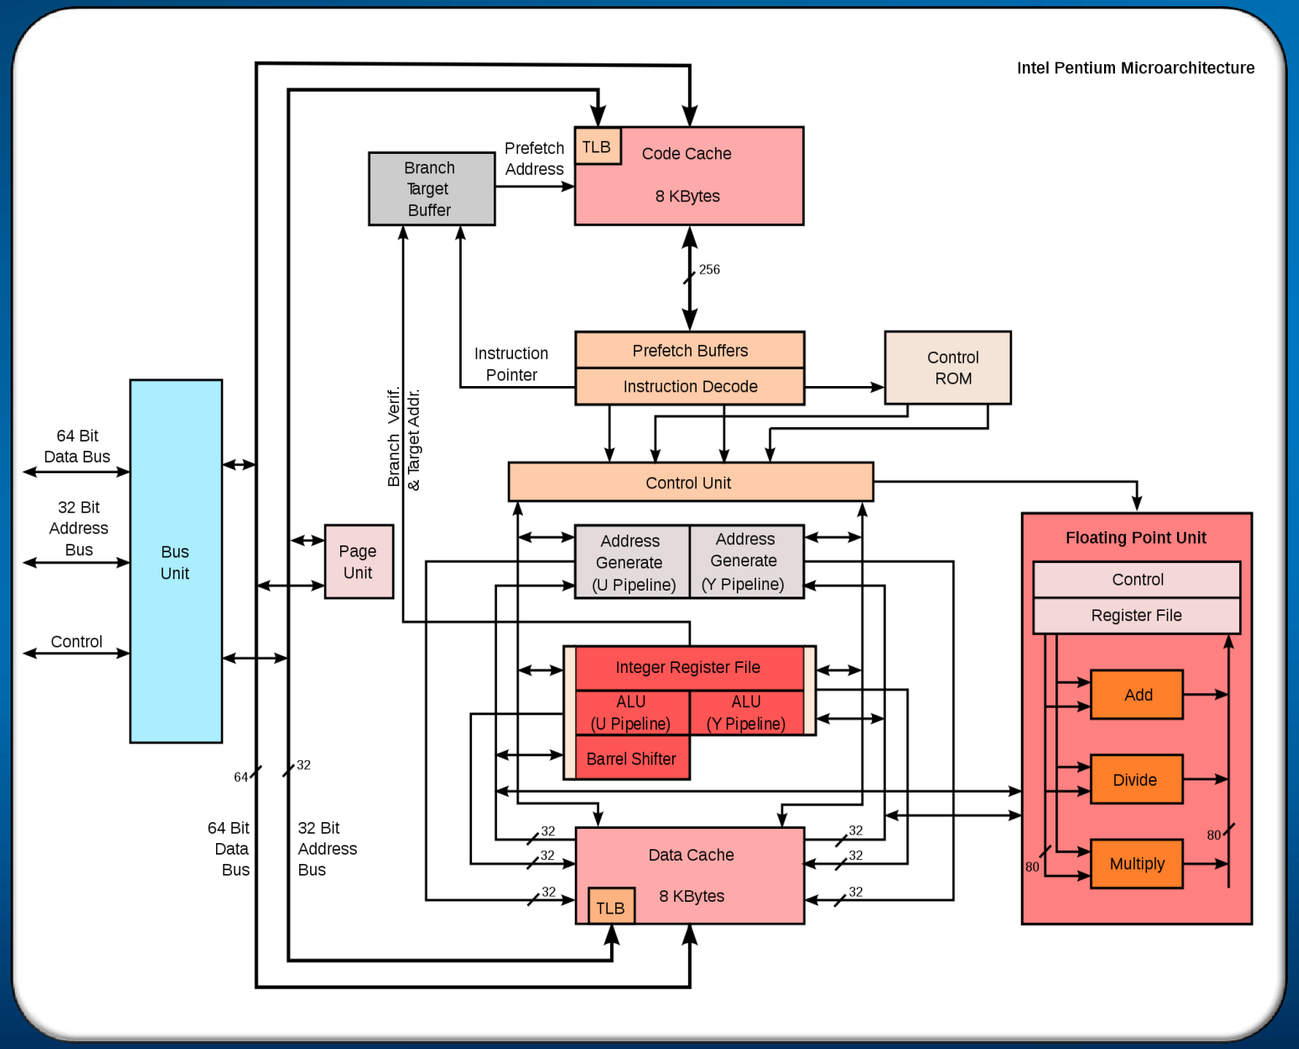

- Legacy Pentium ® core (1994)
- 32-bit CPU with 64-bit data bus
- Superscalar, in-order pipeline architecture with pipelined floating point unit
- Dynamic branch prediction
- Separate code and data caches (8KB)
- Fractional bus operation allowing core frequencies higher than 66MHz FSB

Setting the Design Targets

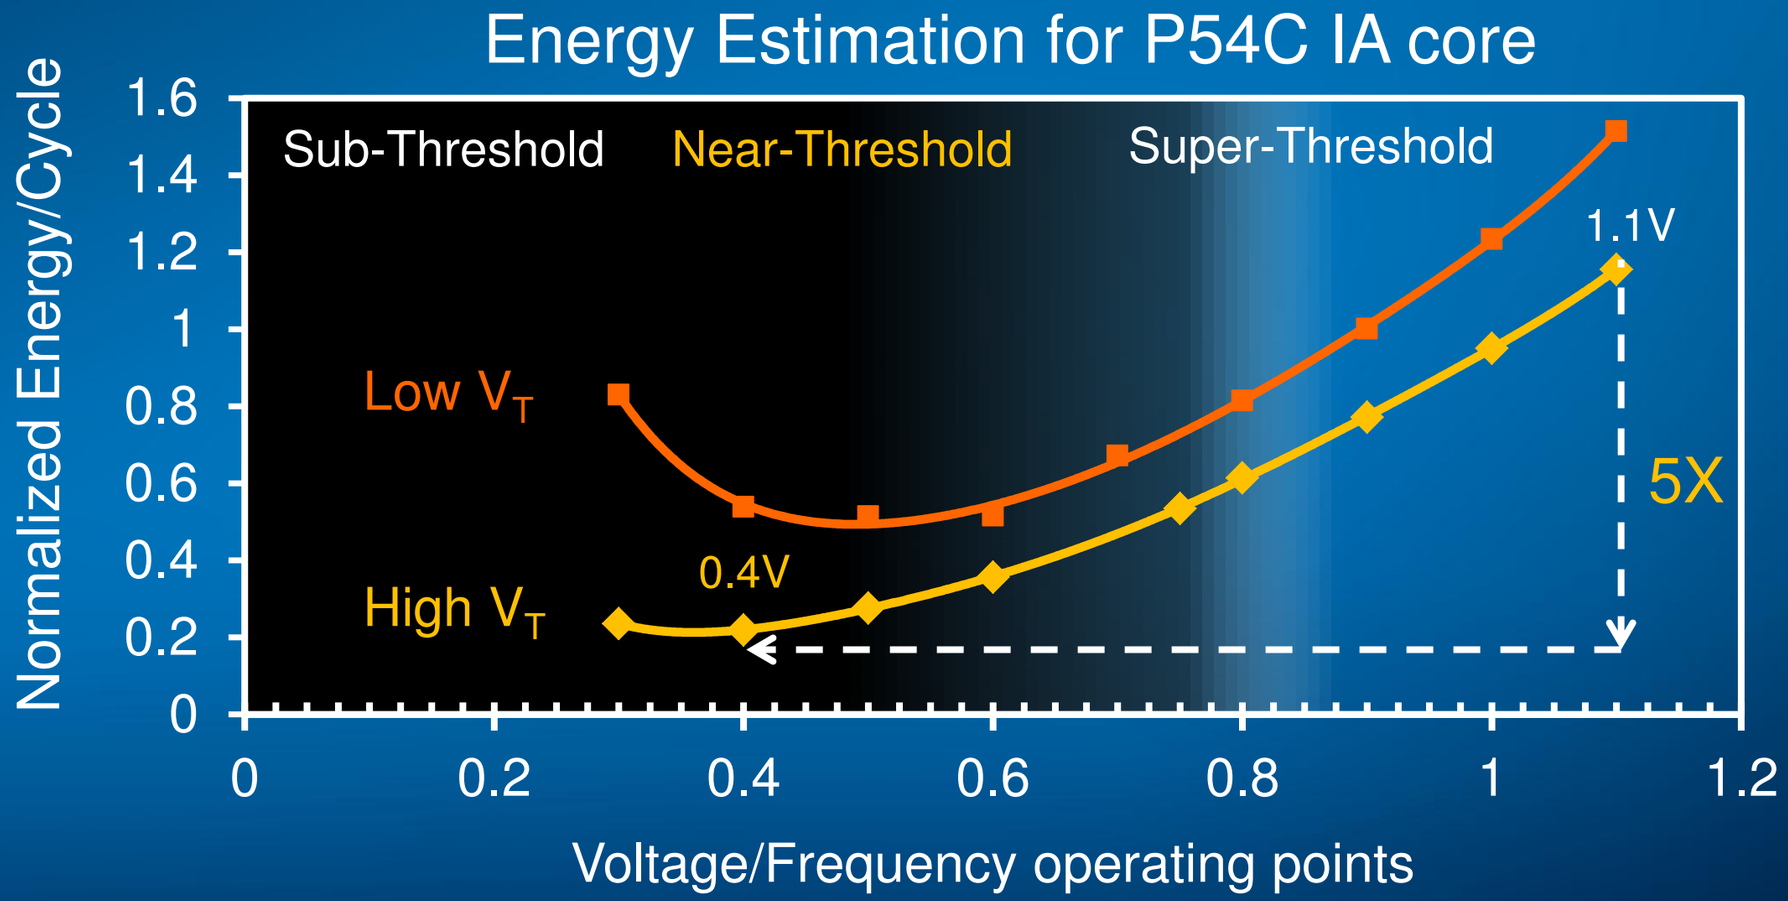

~5X efficiency improvement with aggressive voltage scaling

Claremont Prototype

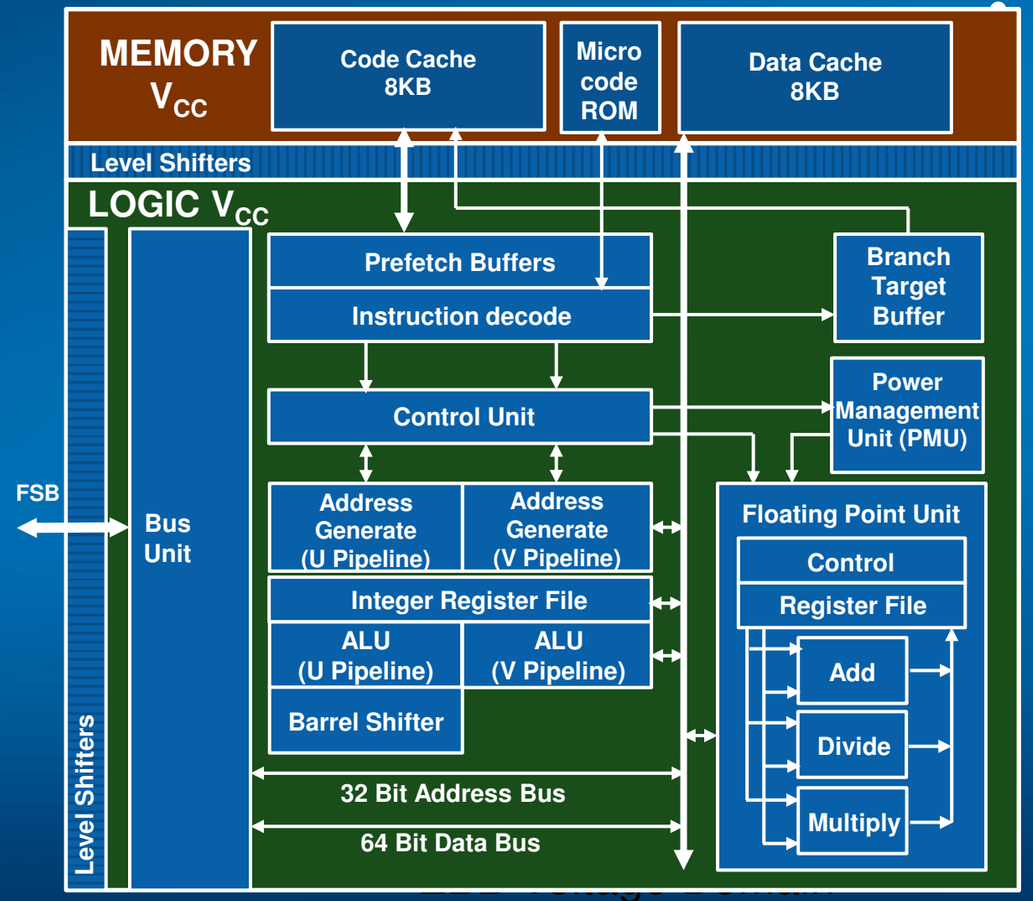

- P54C IA Core in 32nm CMOS
- Aggressive Voltage Scaling
  - Vmin Target: 0.5V (Logic) and 0.55V (RF)
  - Low Voltage, Variation Aware Design
- Ultra Low Power Design
  - Sub-20mW Total Core Power Target at 0.5V
- Wide Dynamic Operating Range
  - Triple Corner Optimization
  - 0.5V/66MHz, 0.75V/333MHz, 1.05V/525MHz

Proactive Power Management

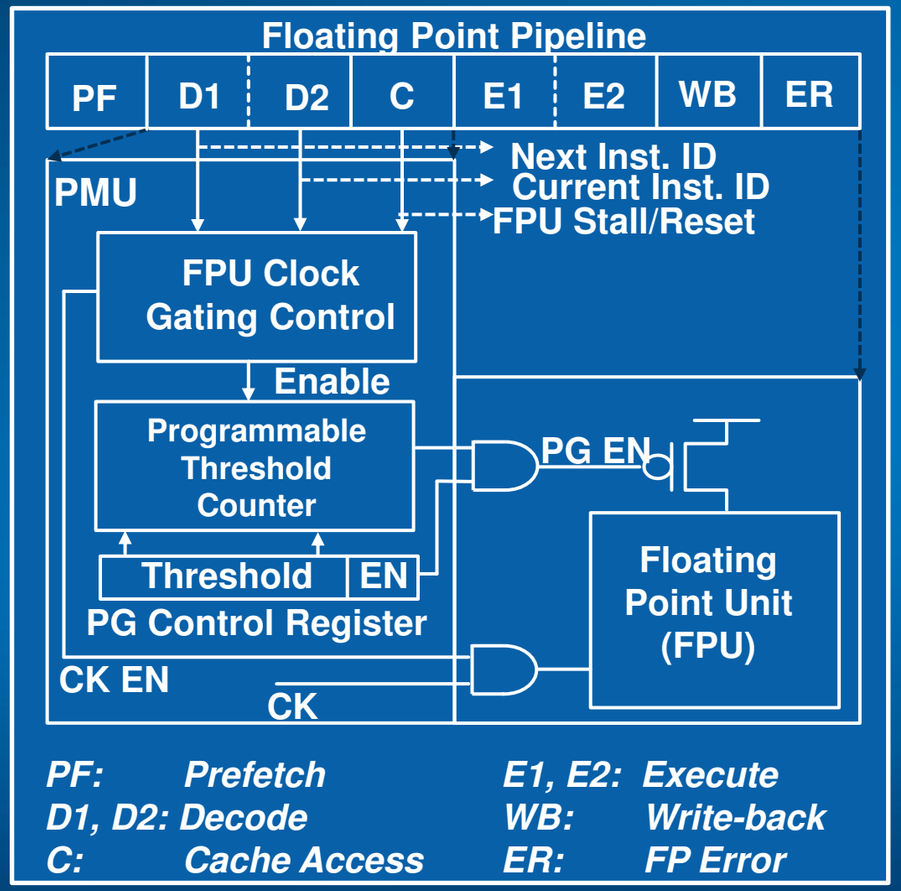

- **Instruction-driven power gating**
  - Dynamically turn-off FPU during idle periods
  - 65% FPU sleep enabled
  - Single cycle wake-up
- **Programmable sleep-threshold**
  - Based on application and operating point
  - Energy saved > Wake-up overheads

Workload-aware, fine-grain power management

Agenda

- Design Challenges
- Claremont Prototype
- Design Strategies and Methodologies
  - NTV Design
  - Wide Dynamic Range Design
- Results and Summary

Near Threshold Voltage Logic

- Variation aware library pruning to ensure reliable NTV operation
- Limited transistor stacks to 3, No wide TG muxes, No contention circuits
- Pruned minimum sized and low drive strength cells
- Sequentials redesign with interruptible and upsized keepers
- 10T single ended transmission gate register file cell topology
- Semi interruptible split output level shifters
- Full swing 3.3V I/O for legacy board compatibility

Low Voltage Timing Convergence

Random/Systematic process variations ➔ Path delay uncertainty ➔ Max/min failures

- Max convergence strategy
  - Setup violations are not fatal; Can be corrected by relaxing PV phase
  - Conservative library pruning ➔ Low variation impact ➔ Better PV to silicon correlation
  - Shallow P54C pipeline (~75 gate stages) ➔ Averaging effect of random variations
- Min convergence strategy
  - Hold violations are fatal; Important to ensure sufficient hold margins
  - Tango charges data/clock path variation; Need to consider sequential hold variation as well..

“Variation aware” timing convergence is essential

Hold Margin Guard Banding

Hold time variation characterization using NOVA/MPP2

Sequential Variants	Hold Time TZST, 20C		
	0.5V, 0σ	0.5V, 5.5σ	0.4V, 5.5σ
✗ 1X Local Clock Inverters	-159 pS	686 pS	9899 pS
✓ 2X Local Clock Inverters	-224 pS	136 pS	786 pS

Variation aware hold margin guard banding for robust sub-0.5V operation

Agenda

- Design Challenges
- Claremont Prototype
- Design Strategies and Methodologies
  - NTV Design
  - Wide Dynamic Range Design
- Results and Summary

Wide Dynamic Range Design Challenge

P1269.4	Critical Path at 0.5V	Critical Path at 1.05V
Device Delay (DD) Contribution	98%	75%
Interconnect Delay (ICD) Contribution	2%	25%
Characteristics	Device dominated data paths become most critical. Effect of ICD is negligible	ICD becomes significant component at higher voltages

The change in critical path distribution across the wide voltage range significantly increases timing optimization efforts to achieve targeted frequency at a given voltage corner

Wide Dynamic Range Design Optimizations

Synthesis Corner Evaluation	Timing Metrics					
	0.75V			0.5V		
	WNS	TNS	Fmax	WNS	TNS	Fmax
✗ 0.75V, 450MHz	-0.18 nS	-165 nS	420 MHz	-1.8 nS	-1950 nS	74 MHz
✓ 0.5V, 86MHz	-0.16 nS	-21 nS	430 MHz	-0.3 nS	-4.6 nS	84 MHz

- Insignificant ICD at ULV ➔ Suboptimal P&R ➔ ICD dominated critical paths at high voltages
  - Prioritized ICD dominated paths at partition level before placement
- Synthesis constraints: Superset of timing constraints required across voltage range

Multi corner constraints for single corner synthesis

Multi-Voltage Clocking and Skew Management

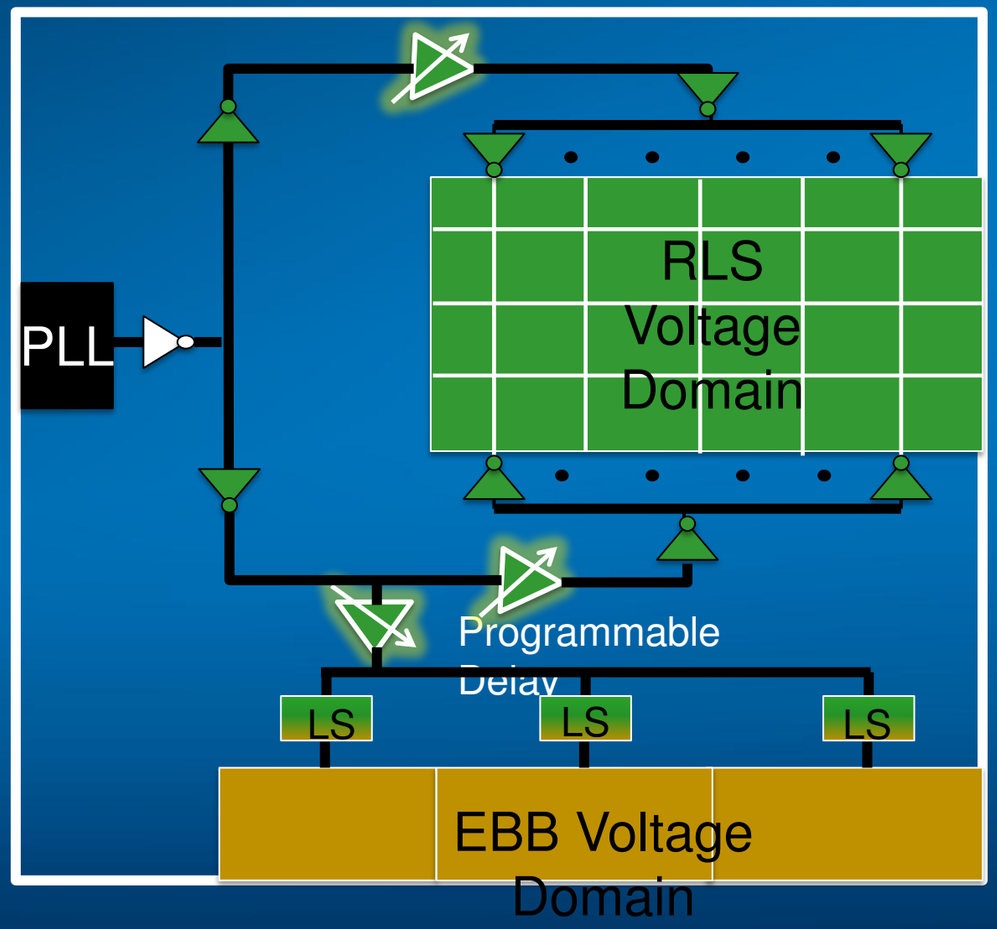

- Core logic (RLS) & memories (EBB) operate on independent power supply
- Level shifters in clock distribution network
- Inter-block skew is voltage dependent
- Programmable delay buffers for the skew management, configured via scan
- 1.5-2X skew reduction across different RLS, EBB voltage combinations

Enhanced PV Methodology

- Modified intra-block & inter-block skew function in Tango for multi-voltage timing analysis and roll-ups
- Incorporated block Voltage Mapping (VM) table and Voltage Dependent Skew (VDS) table in Tango environment
- Skew is computed based on operating voltage of launching and sampling block using entries in VM & VDS tables

Agenda

- Design Challenges
- Claremont Prototype
- Design Strategies and Methodologies
  - NTV Design
  - Wide Dynamic Range Design
- Results and Summary

Claremont Die Micrograph and Test Setup

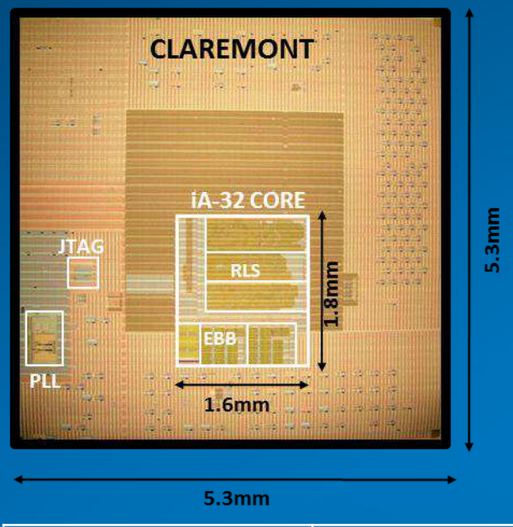

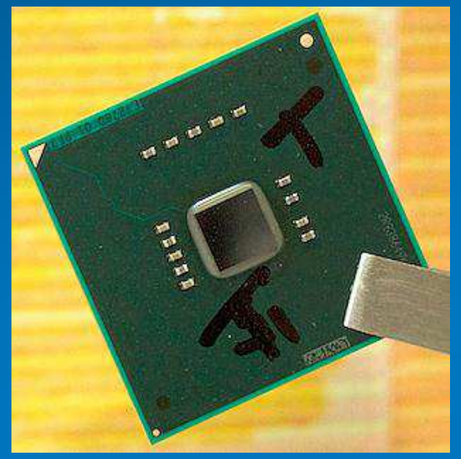

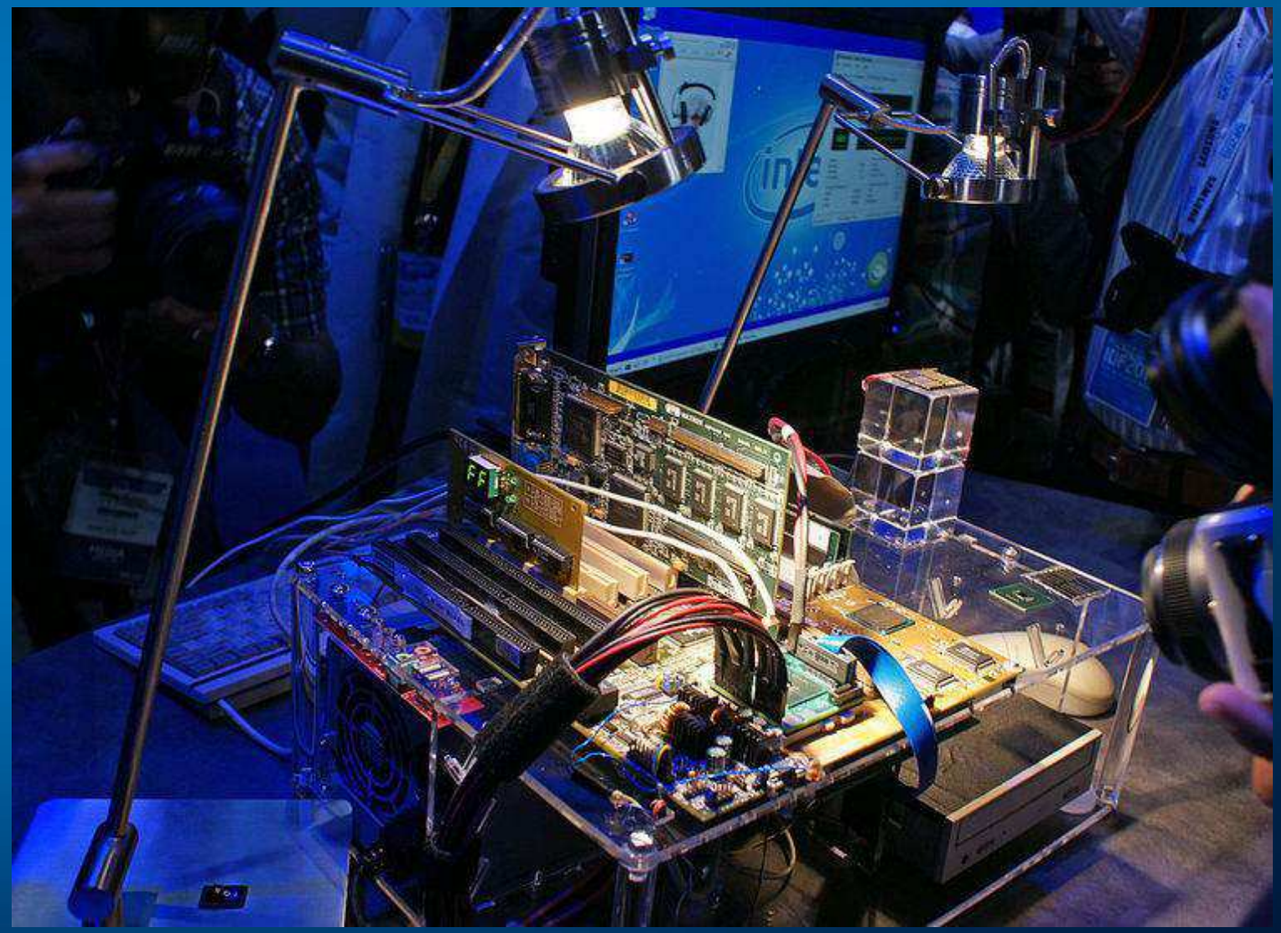

Technology	32nm High-K Metal Gate Technology
Interconnect	1 Poly, 9 Metal (Cu)
Transistors	6 Million (core & EBBs)
Core Area	1.96mm^2
Signals	168
Package	951 Pins FCBGA11

Claremont Test Challenges

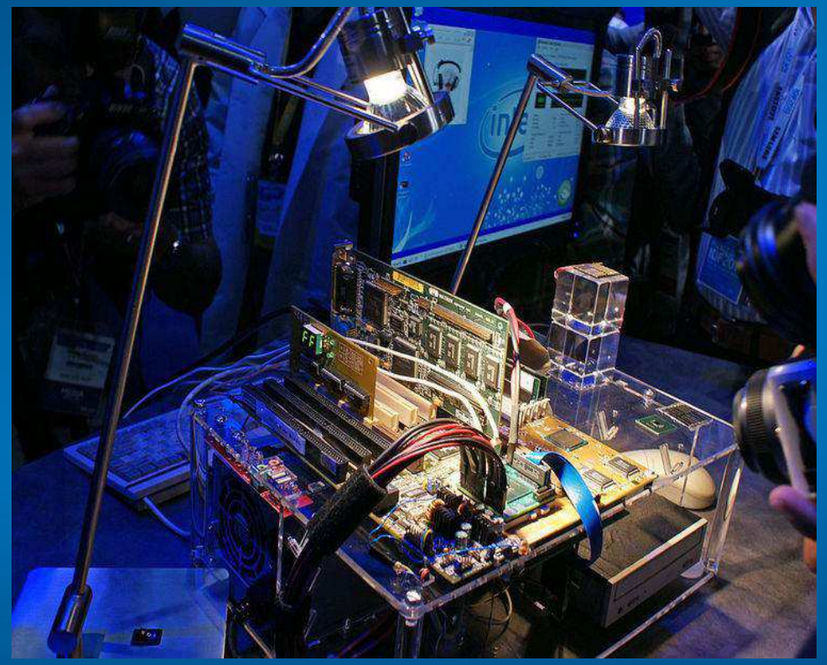

- 18 year old motherboard
  - Age variation
- Most peripherals fail below 15Mhz
- Front side bus spec timing and voltage challenges
- Lack of uBreak points & advanced debug features

Power/Performance Characteristics

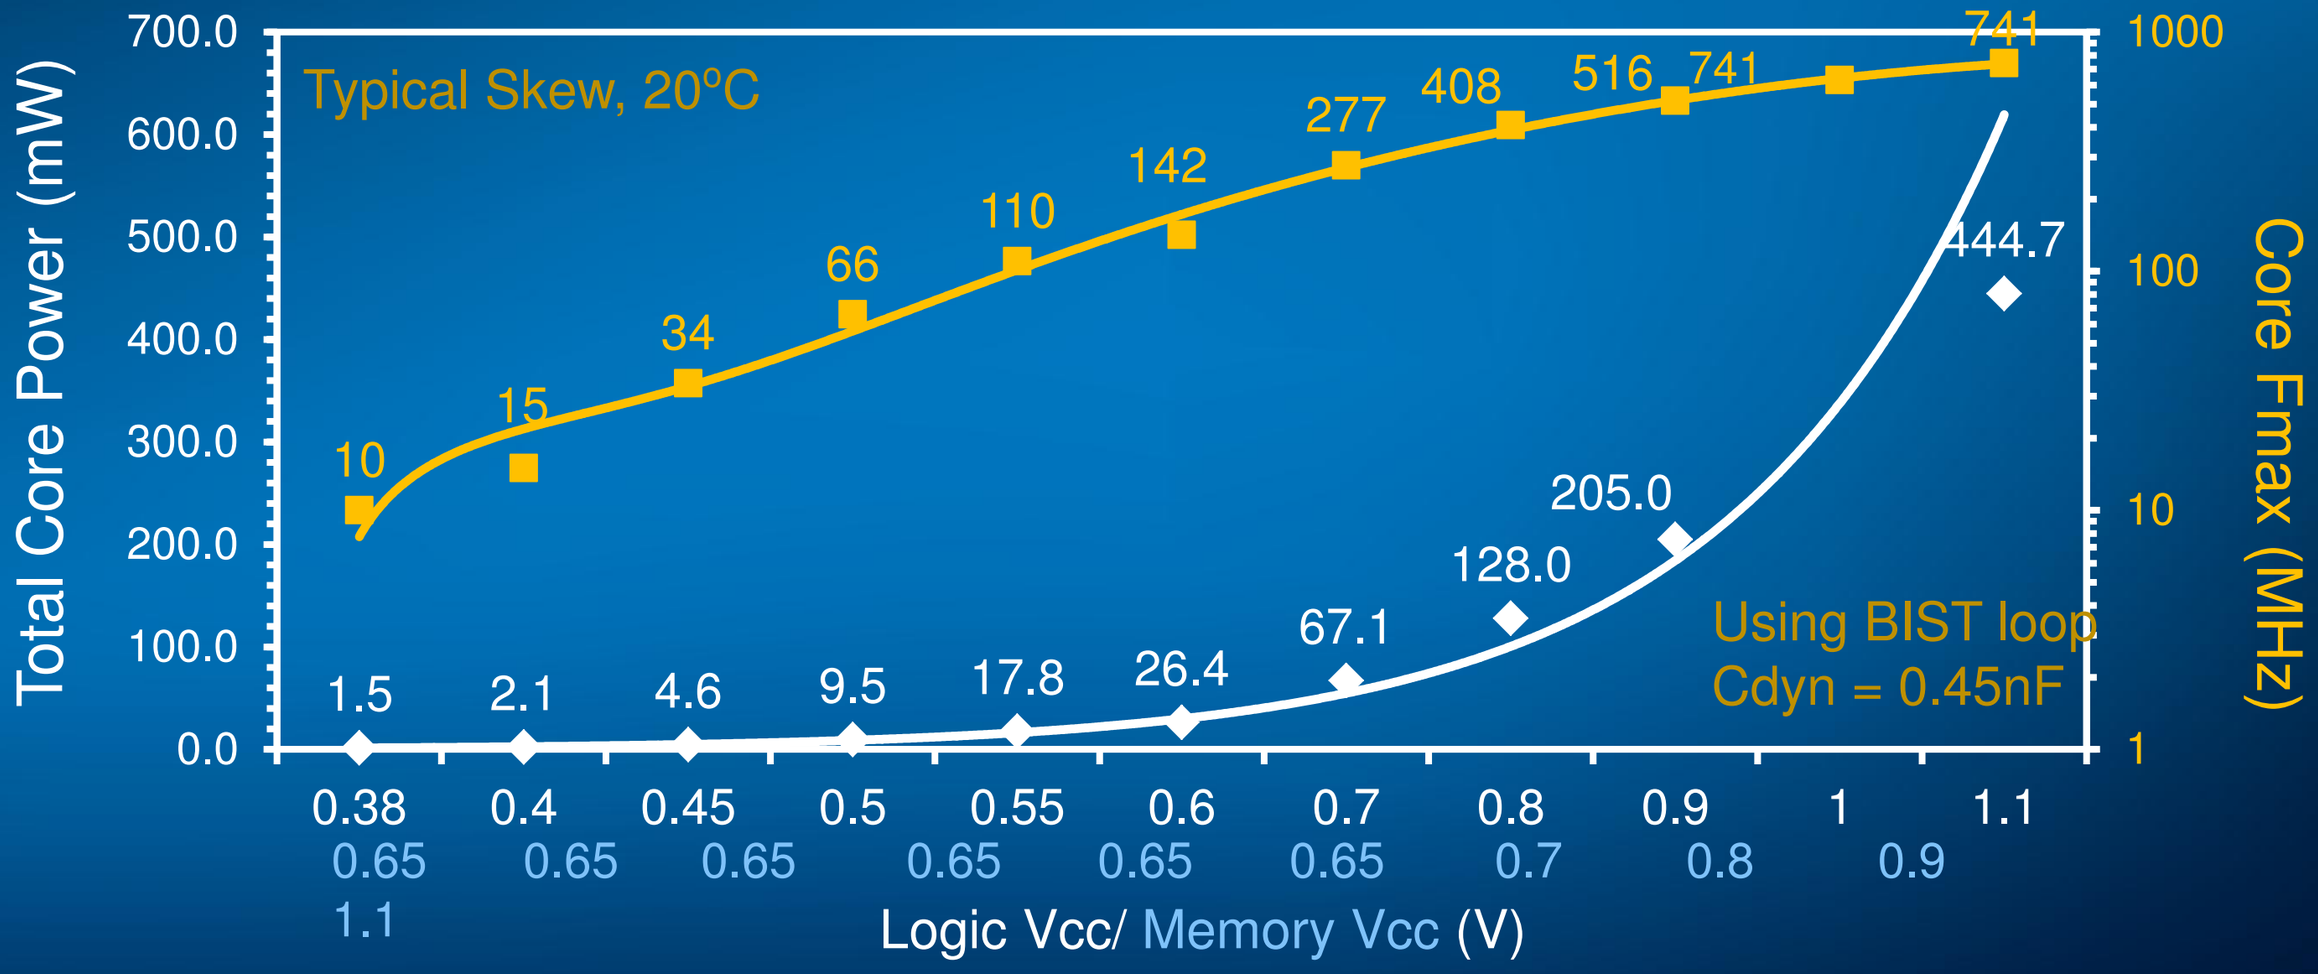

Wide Dynamic Range 1.1V/741MHz/445mW to 380mV/10MHz/1.5mW

Total Power Breakdown

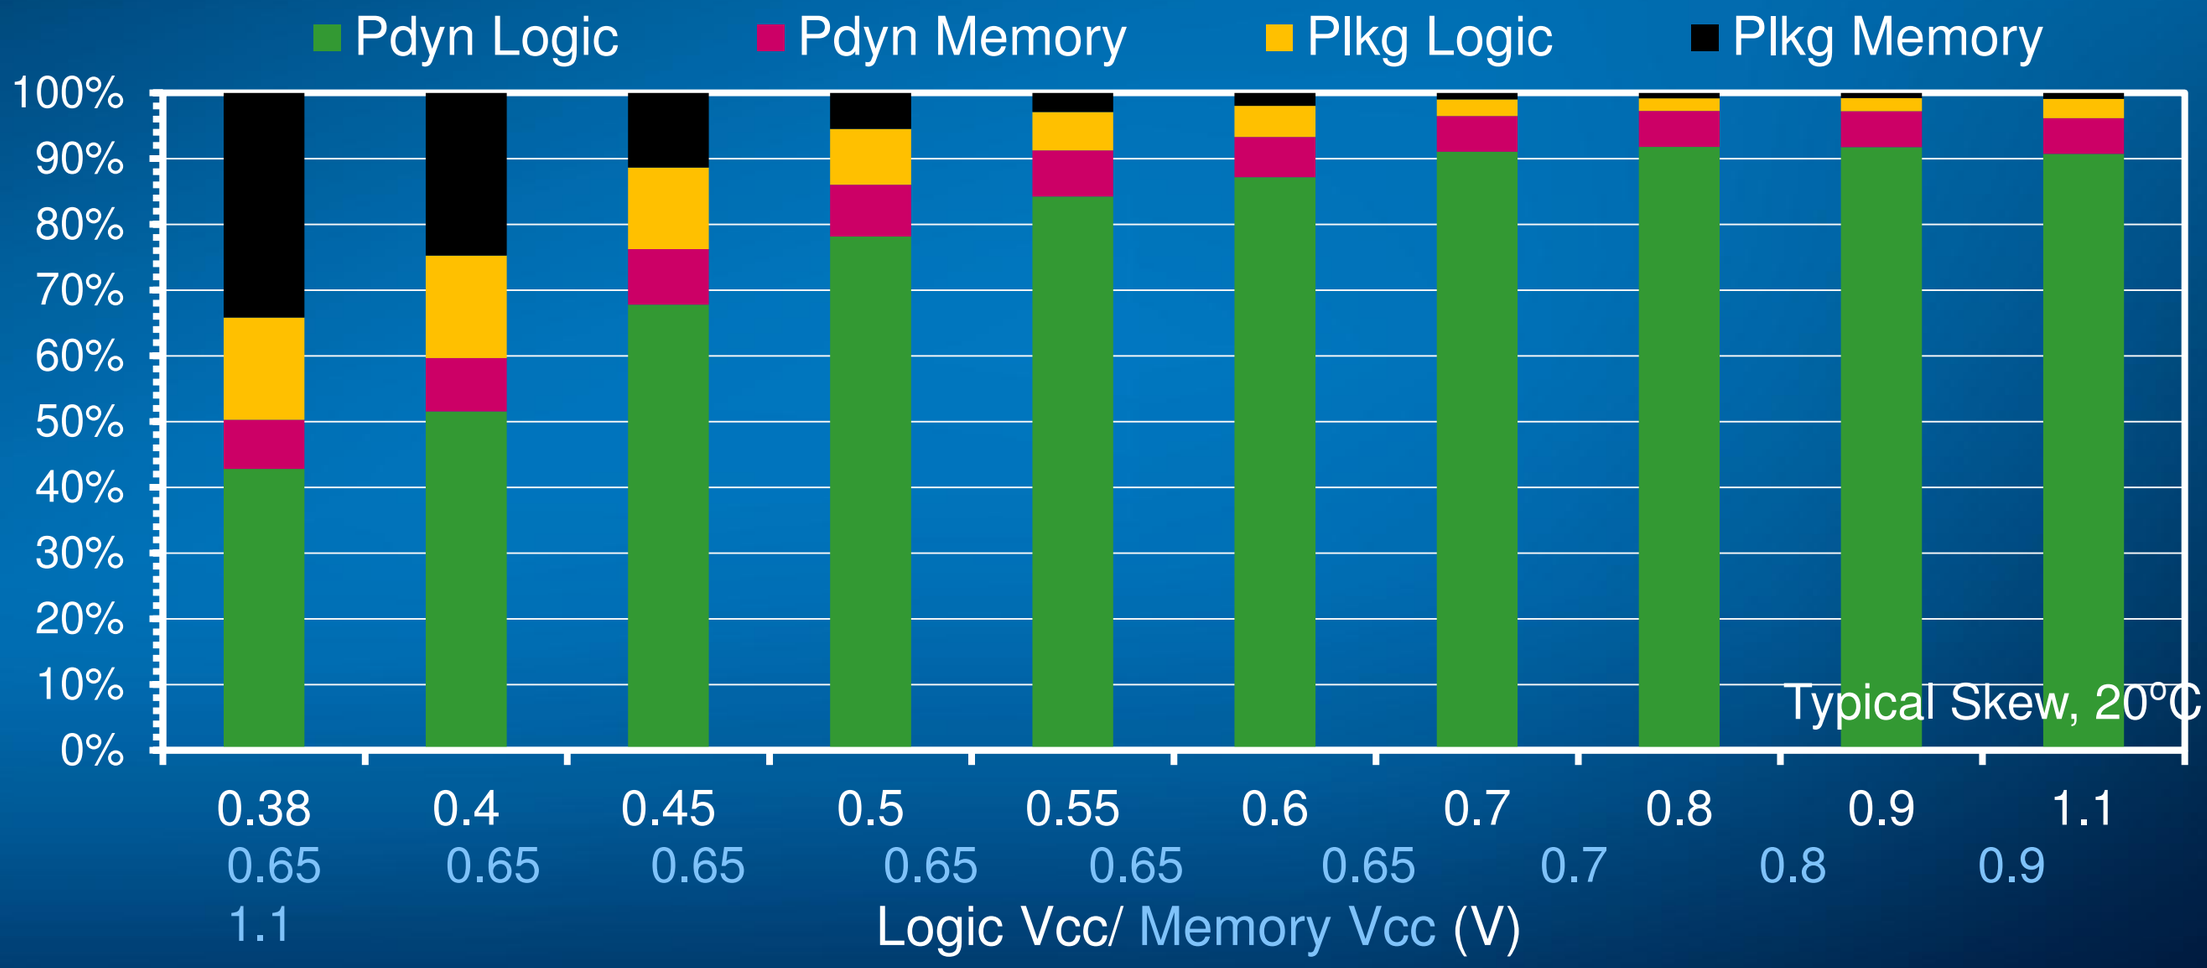

Leakage power scales from 3% @1.1V to 50% @ 0.38V

Energy/Cycle: Typical Skew

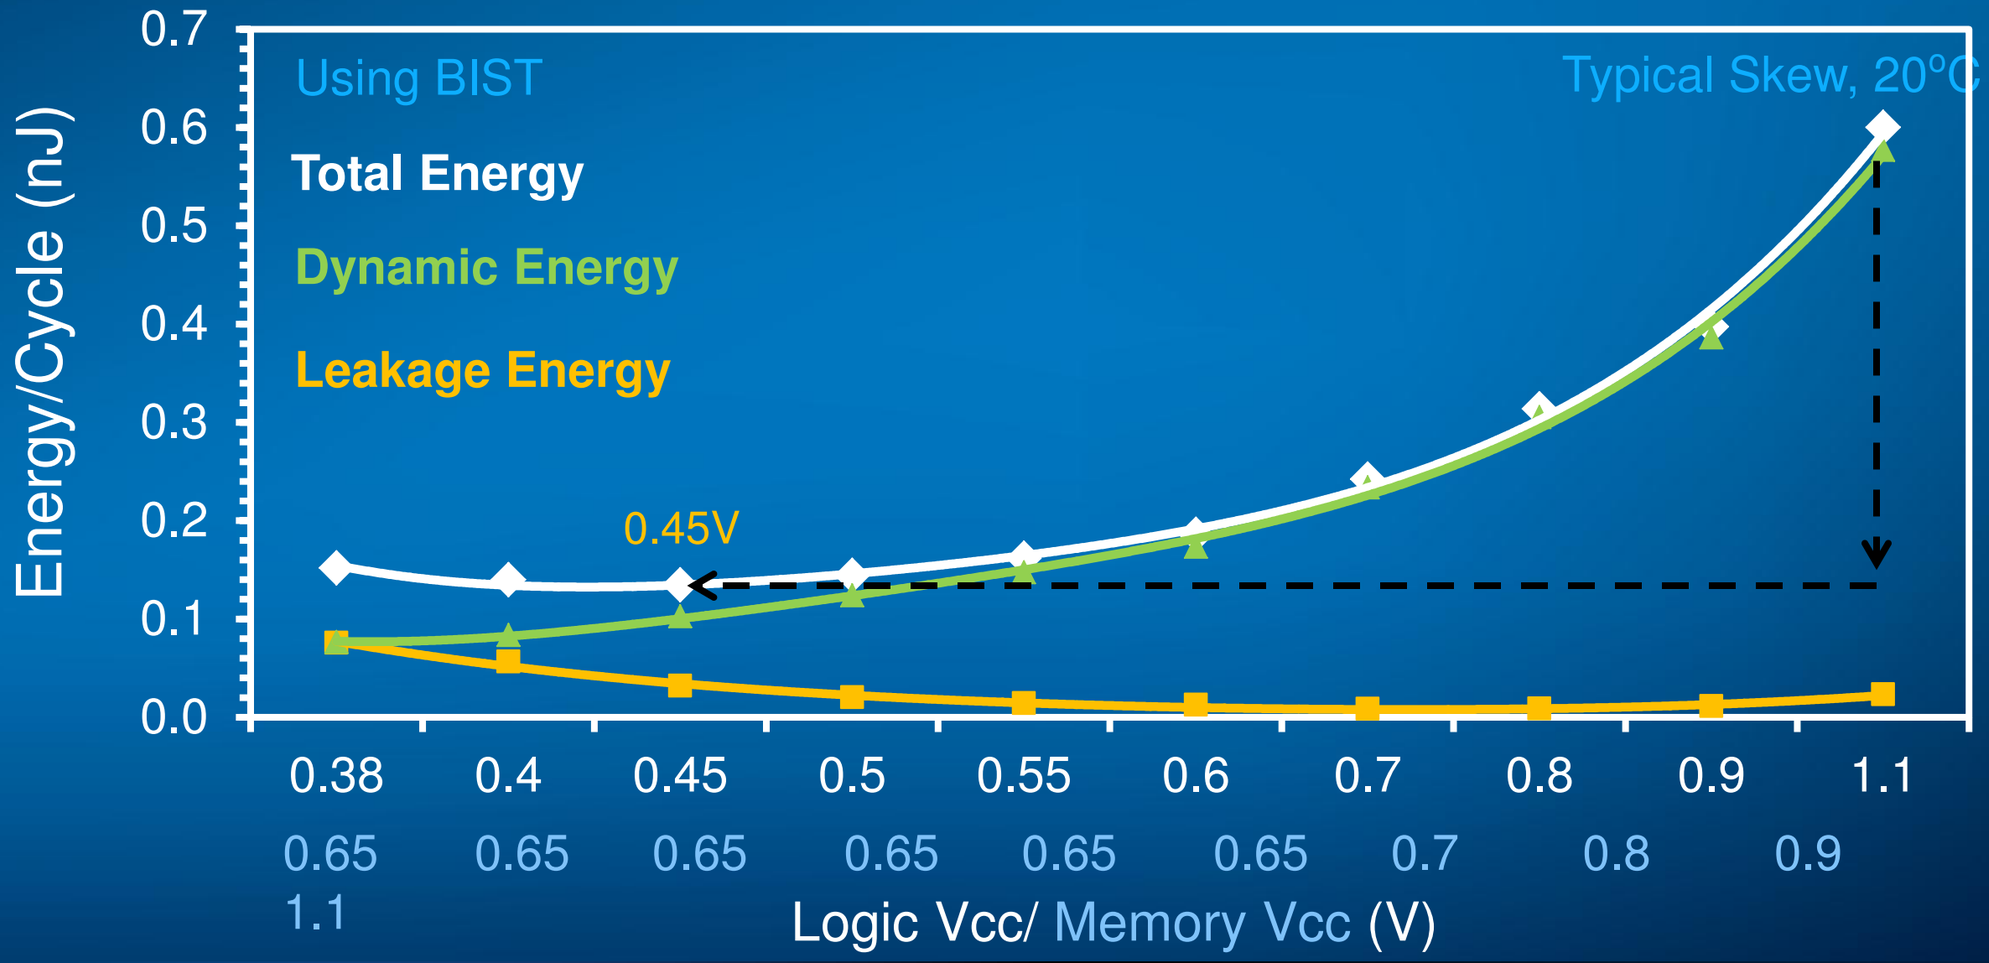

4.5X Energy reduction from Vmax: 135pJ/cycle at 450mV

Temperature Dependence: Typical Skew

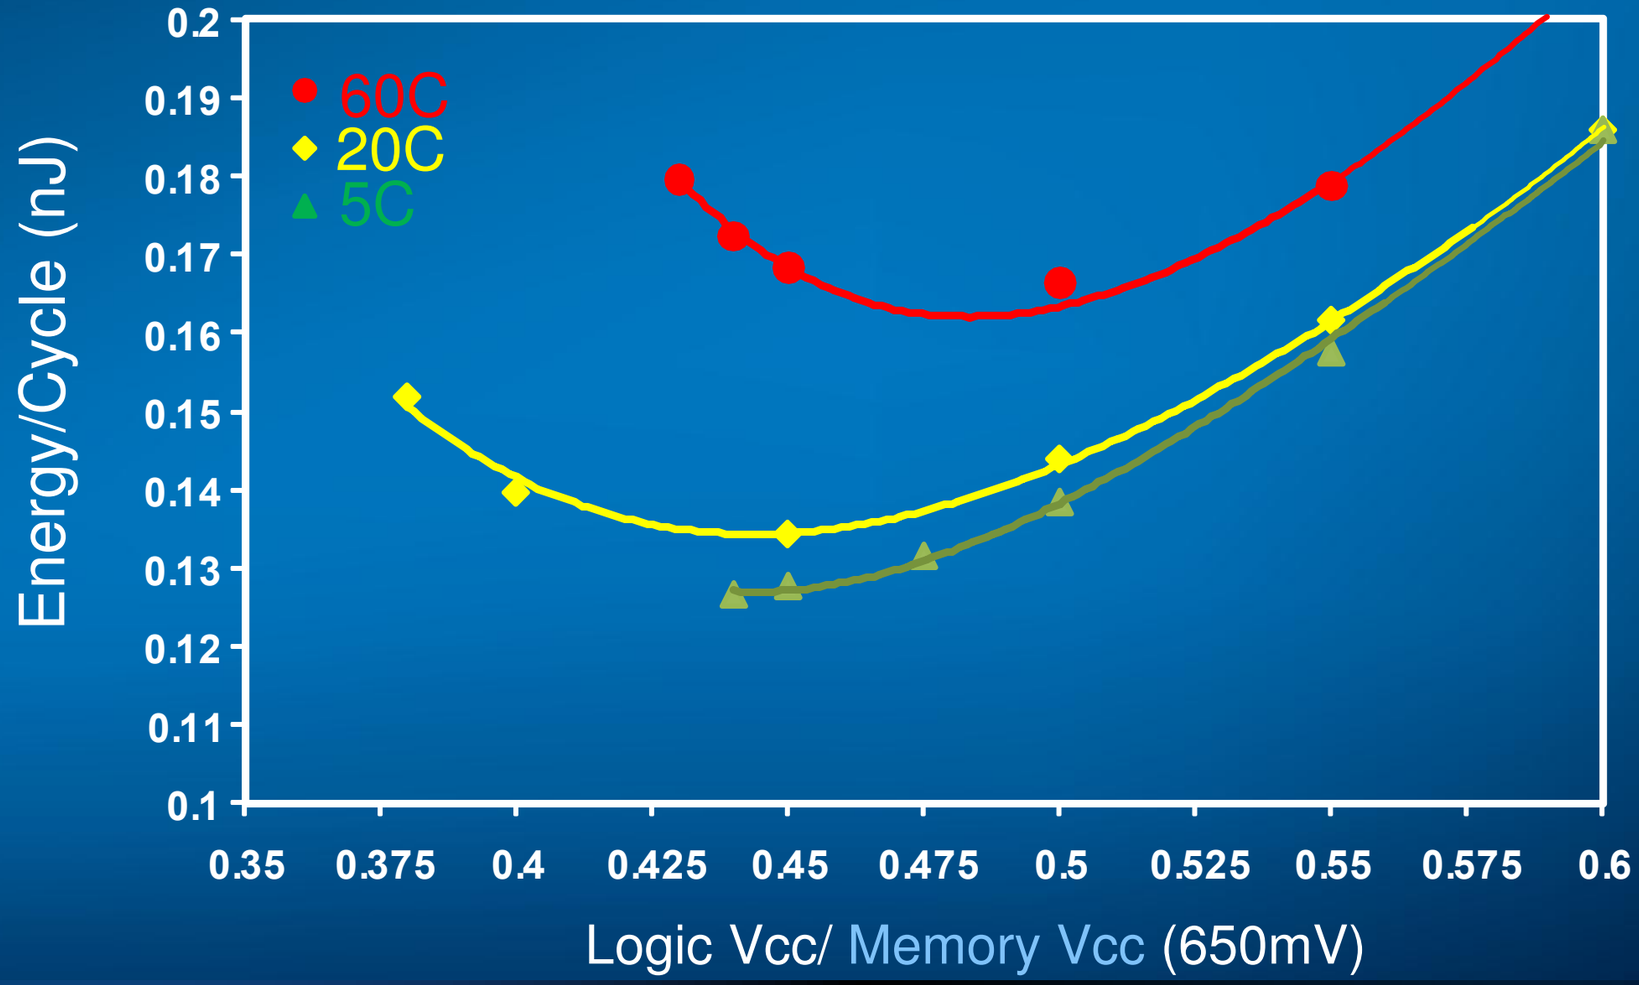

25% increase from 5C to 60C

Energy/Cycle: Fast Skew

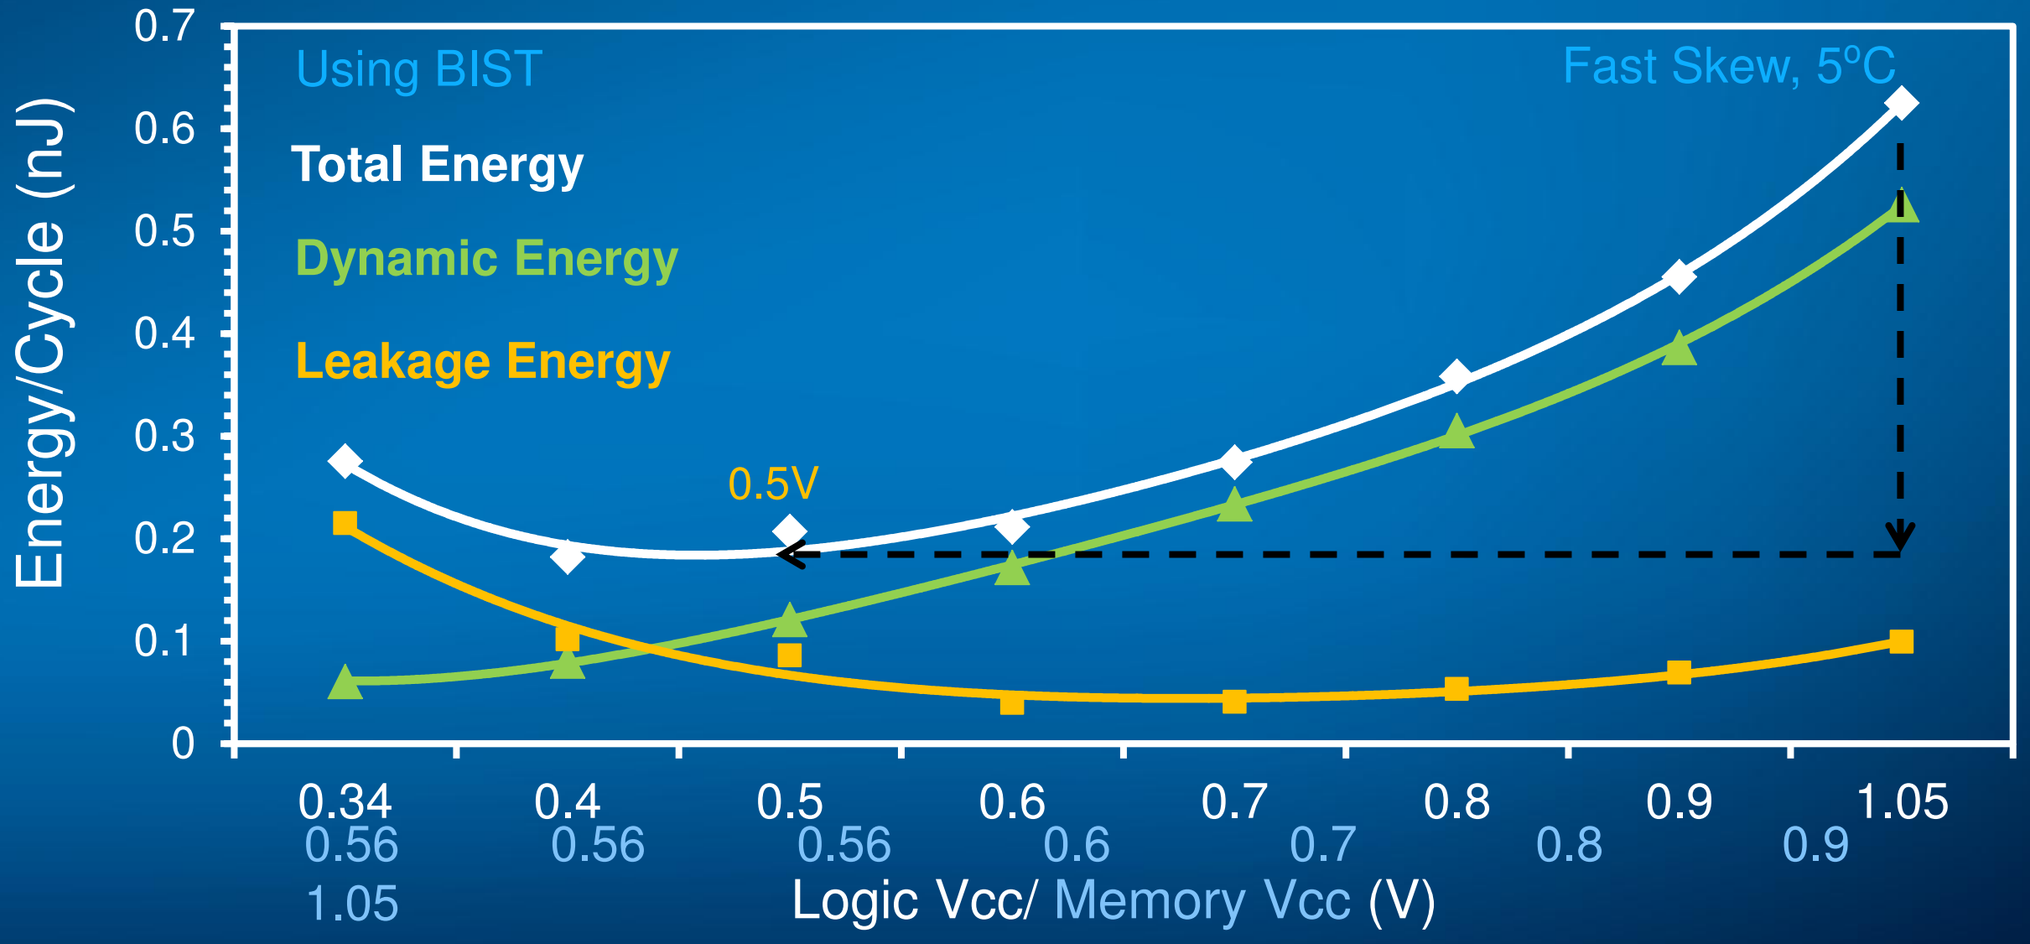

3X Energy reduction from Vmax: 200pJ/cycle at 500mV

Area Penalty per Technique

Technique	Increase from Audit DB
Modified Sequentials	27%
0.5V 66MHz Target	24%
Complex cells Pruning	10%
3 RLS blocks vs. 1	8%
Min Z	5%
Additional Min fixing	2%

- **Area overhead is a non-linear function of Vccmin improvement.**
- **Incremental Vccmin improvement is more practical and will have a lower penalty**

Claremont: Industry's First NTV IA Core

380mV

Core Logic Vmin

1.5mW

Total Core Power at Vmin

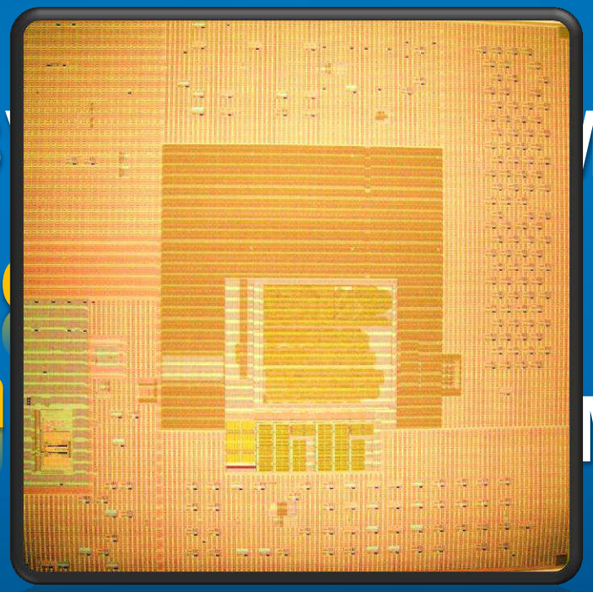

293X

Total Power Reduction from Vmax to Vmin

4.5X

Total Energy Savings from Vmax to Vopt

A0 silicon booting Multiple O/S

Measured using BIST

Conclusion

- NTV: A Promising Technology for Energy Efficient Computing
  - Beneficial for ultra-low power IA with modest performance demands
  - Energy efficient SOCs, Graphics, Sensor hubs, Many-core CPUs, Exascale…
- Claremont Demonstrated “Reliable” NTV Operation, Enabled by
  - Novel circuit design techniques for logic, sequentials and memories
  - Variation aware design convergence strategies
- Next Steps:
  - Low overhead NTV circuits, ULV standard cell libraries
  - CAD Methodologies for Low Vcc & WDR designs : SSTA, Multipoint optimization
  - Device – Circuits – Architecture co-design for Near Threshold Computing

Q&A

Backup

Near Threshold Voltage Logic

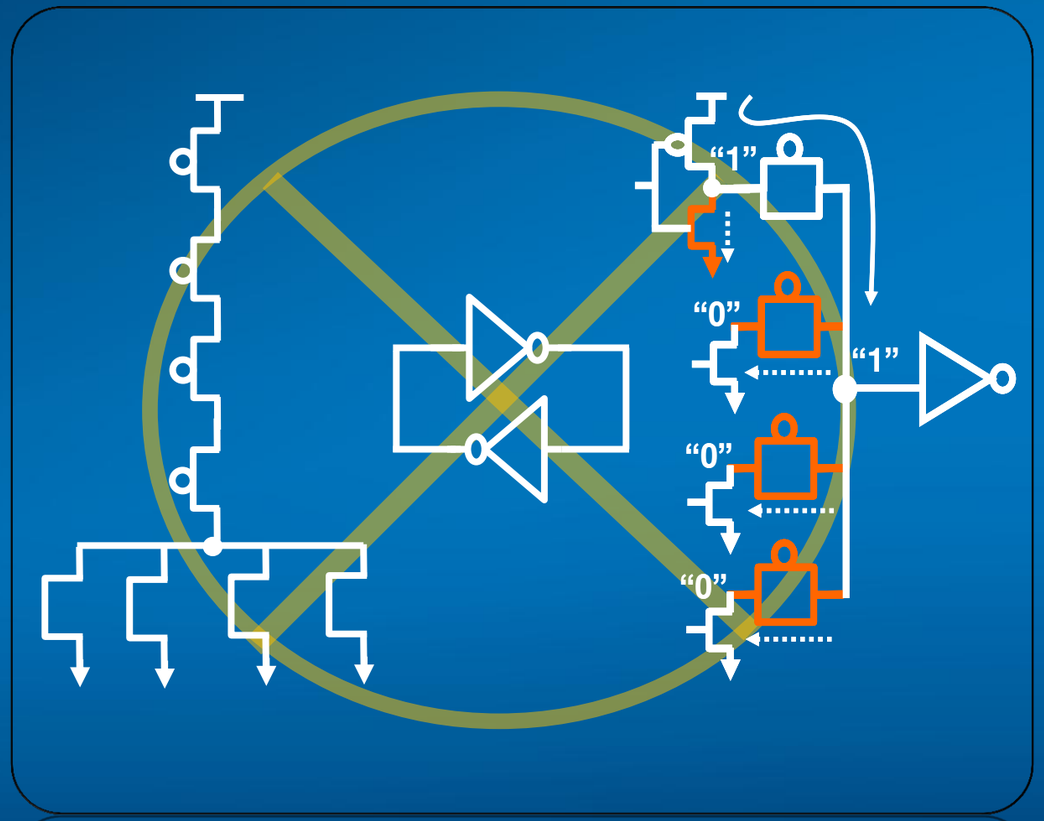

- Variation aware library pruning to ensure reliable NTV operation
- Limited transistor stacks to 3, No wide TG muxes, No contention circuits
- Pruned minimum sized and low drive strength cells: Minimum Z allowed is 2X process Z_{min}
- Only 40% of combinational standard cells in the library used in the design

Re-characterized constrained standard cell library at 0.5V NTV corner

Near Threshold Voltage Sequentials

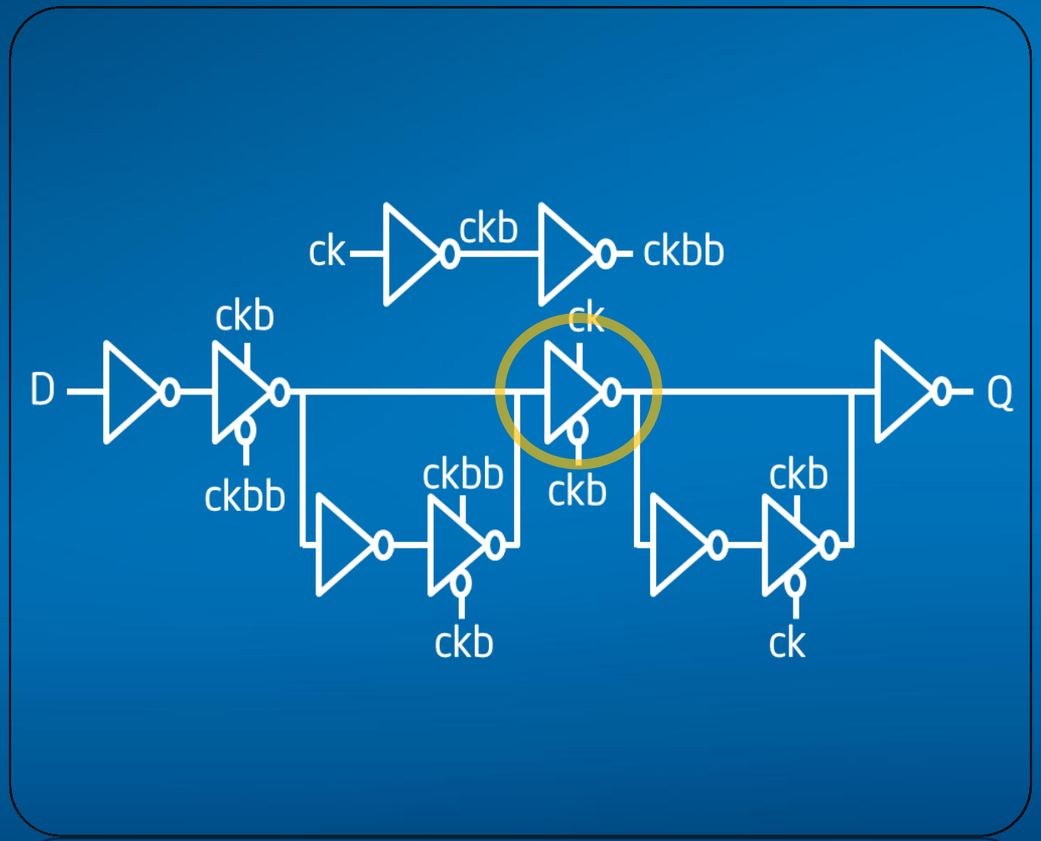

- Modified-C^{2}MOS topology with interruptible and upsized keepers
- Slave TG in conventional flip-flop replaced by clocked inverter, eliminating risk of write-back failures due to charge sharing
- Keepers upsized by 2X to achieve retention Vmin of 0.5V (RSSS, 5.5σ, -25^{o}C, 1.1MΩ Rg)
- Designed 13 custom sequential flavors and re-characterized at 0.5V

Robust, write-back free sequentials with 70mV retention Vmin improvement

Low Voltage RFs

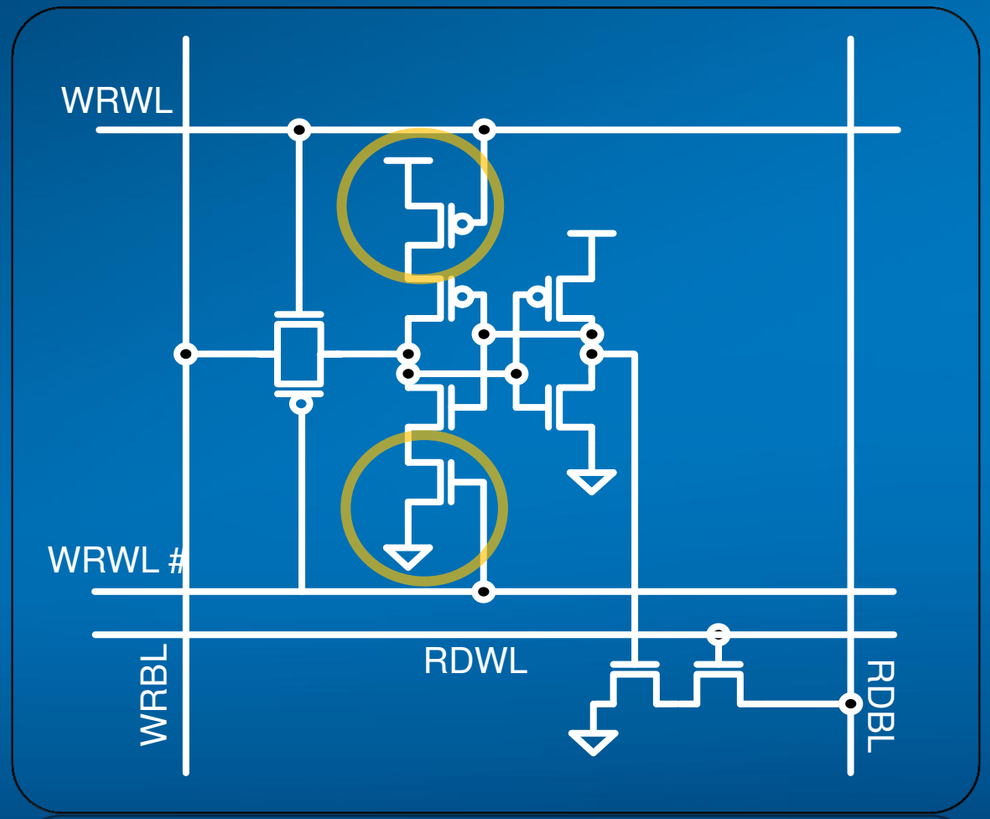

- 10-T single ended transmission gate (SETG) latch topology
- Fully interruptible bit cell for contention free writes
- Retention limited cell, upsized (from 3-Track to 5-Track) to achieve retention Vmin of 550mV (RSSS, 5.9σ, -25ºC, 1.1MΩ Rg)
- Programmable keepers (3 vs. 4 stack) during read

Near Threshold Voltage Level Shifters

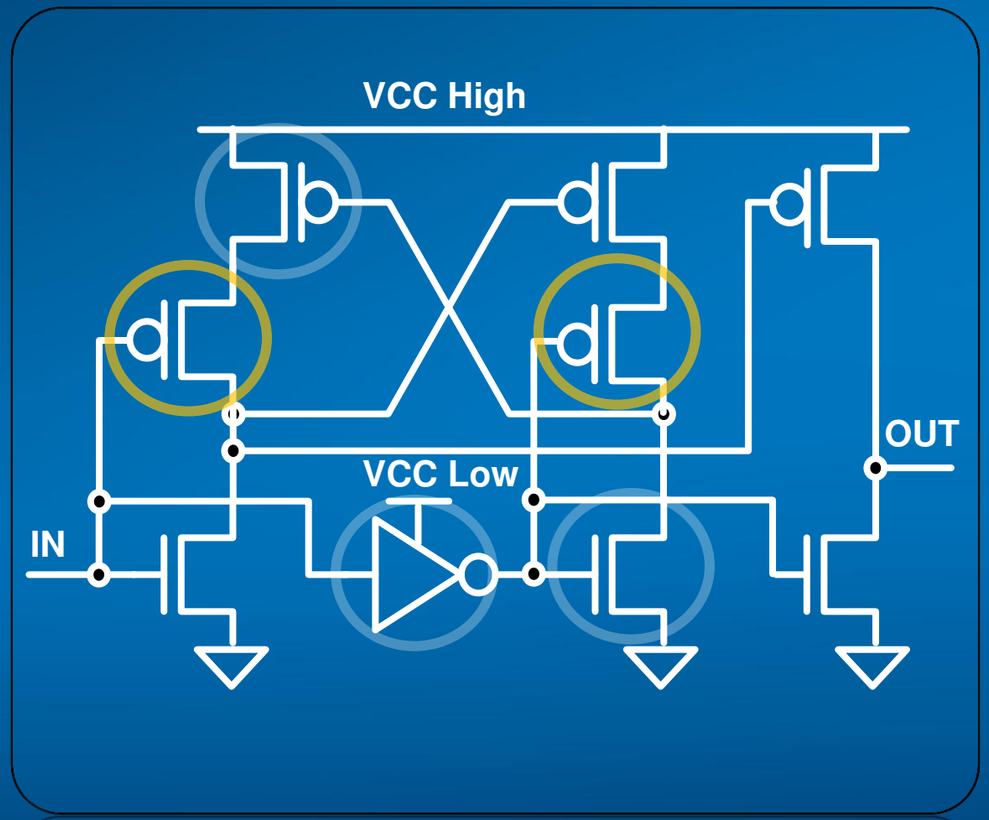

- Semi-interruptible split output level shifter topology
- Interruptible PMOS reduces contention, split output decouples delay path from latch stage
- Asymmetric critical path based sizing
- Two-stage level shifters between RLS & EBB with intermediate reference voltage
- Wide range level translation: sub-0.5V to 1.1V

60% performance improvement over single stage, symmetric level shifters

Package and Evaluation Board

951 Pin FCBGA Package

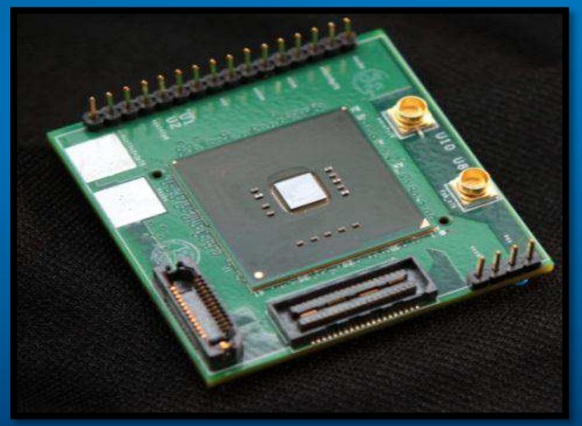

Custom Interposer

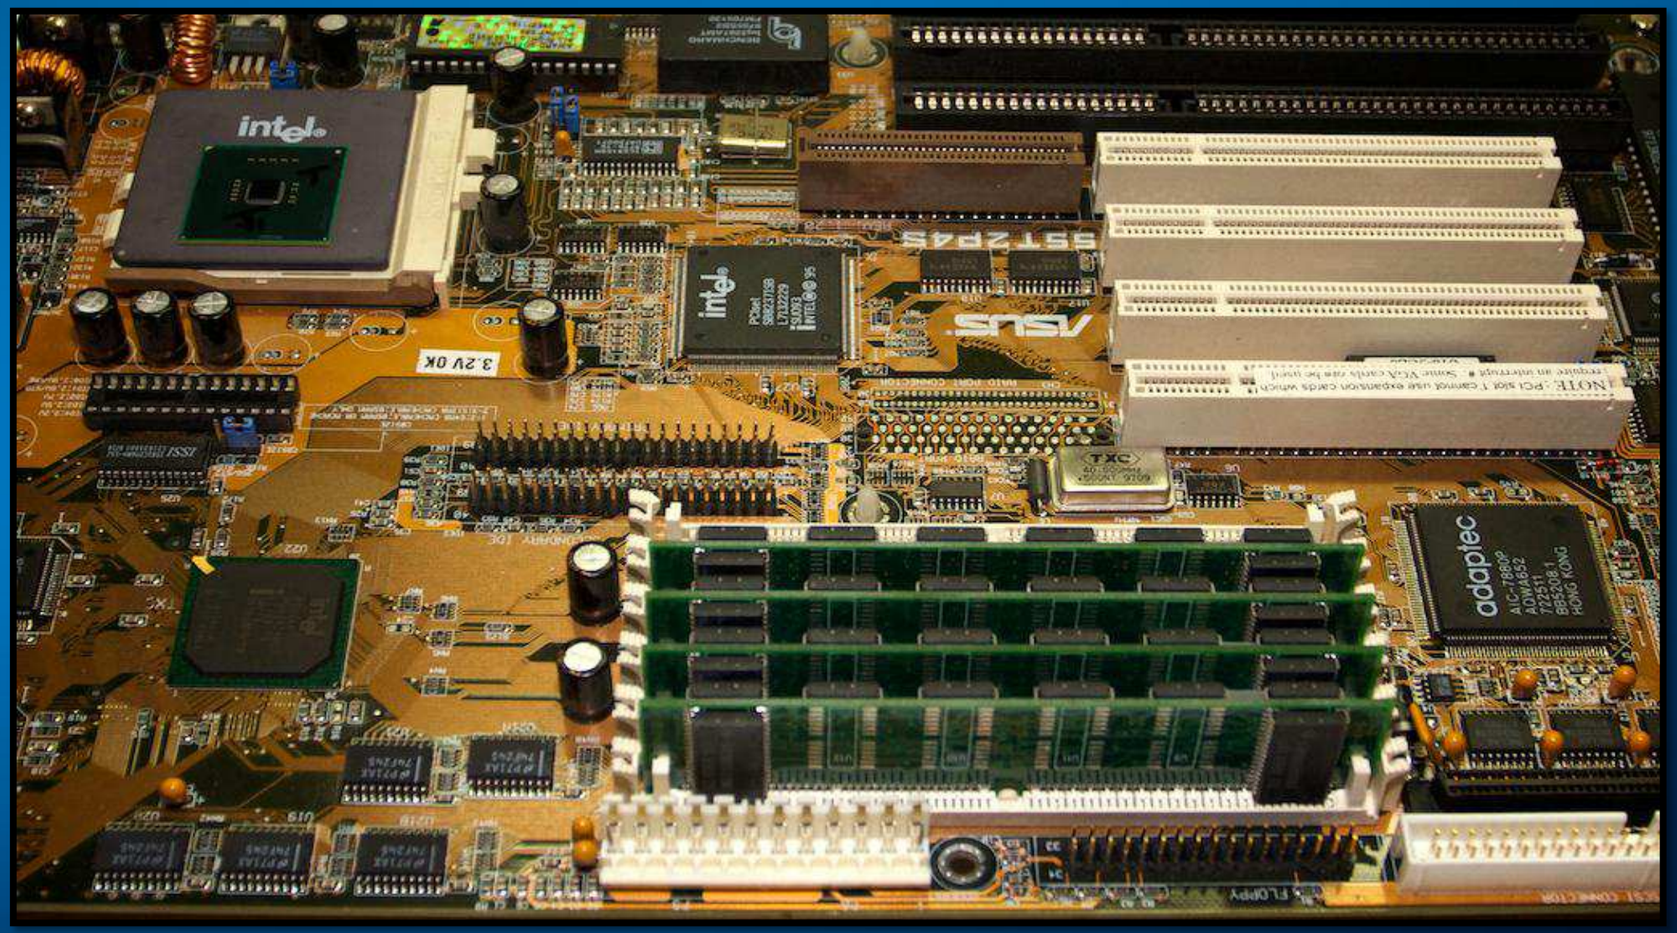

Legacy Pentium™ Socket-7 Motherboard

Legacy Benchmarks

Whetstone (inc. FPU):

- MWIPS MFLOPS VAXMIPS MWIPS-DP
- Pentium 100 66.2 16.8 97.8 66.2 1994
- Pentium 120 79.5 20.2 118 81.6 1995
- Pentium 133 88.3 22.4 130 90.8 1995

Dhrystone: (no FPU)

Dhry1 Dhry1 Dhry2 Dhry2
Opt NoOpt Opt NoOpt

- 80486 DX2 66 45.1 12.0 35.3 12.4
- Pentium 75 112 19.3 87.1 18.9
- Pentium 100 169 31.8 122 32.2
- Pentium 133 239 38.3 181 39.0
- Pentium 166 270 43.6 189 43.9

Linpack: (FPU heavy)
Opt No opt

- 80486 DX2 66 2.63 1.74
- Pentium 75 7.56 4.04
- Pentium 100 12.07 5.40
- Pentium 133 17.05 5.60
- Pentium 166 19.89 6.86