# Implementation and Productization of a 4th Generation 1.8GHz Dual-Core SPARC V9 Microprocessor

Anand Dixit, Jason Hart, Swee Yew Choe,
Lik Cheng, Chipai Chou, Kenneth Ho,
Jesse Hsu, Kyung Lee, John Wu

Sun Microsystems, Sunnyvale, CA, USA

# Overview

- Architecture and Functionality Enhancements
- Physical Implementation Scope
- Library and Design Structure
- Power Management
- Fullchip Integration and Floorplan
- Clocking Design and Analysis
- IO Modification
- Memory Design Changes

# New Architecture Features

- Dual 3rd Generation Enhanced Cores
- Doubled Instruction cache size
- Shared on-chip 2MB Level-2 cache
- Shared external 32MB Level-3 cache
- On-chip Level-3 cache tags
- Optimized System Interface
- Optimized Memory Interface

# Memory Hierarchy of Processor

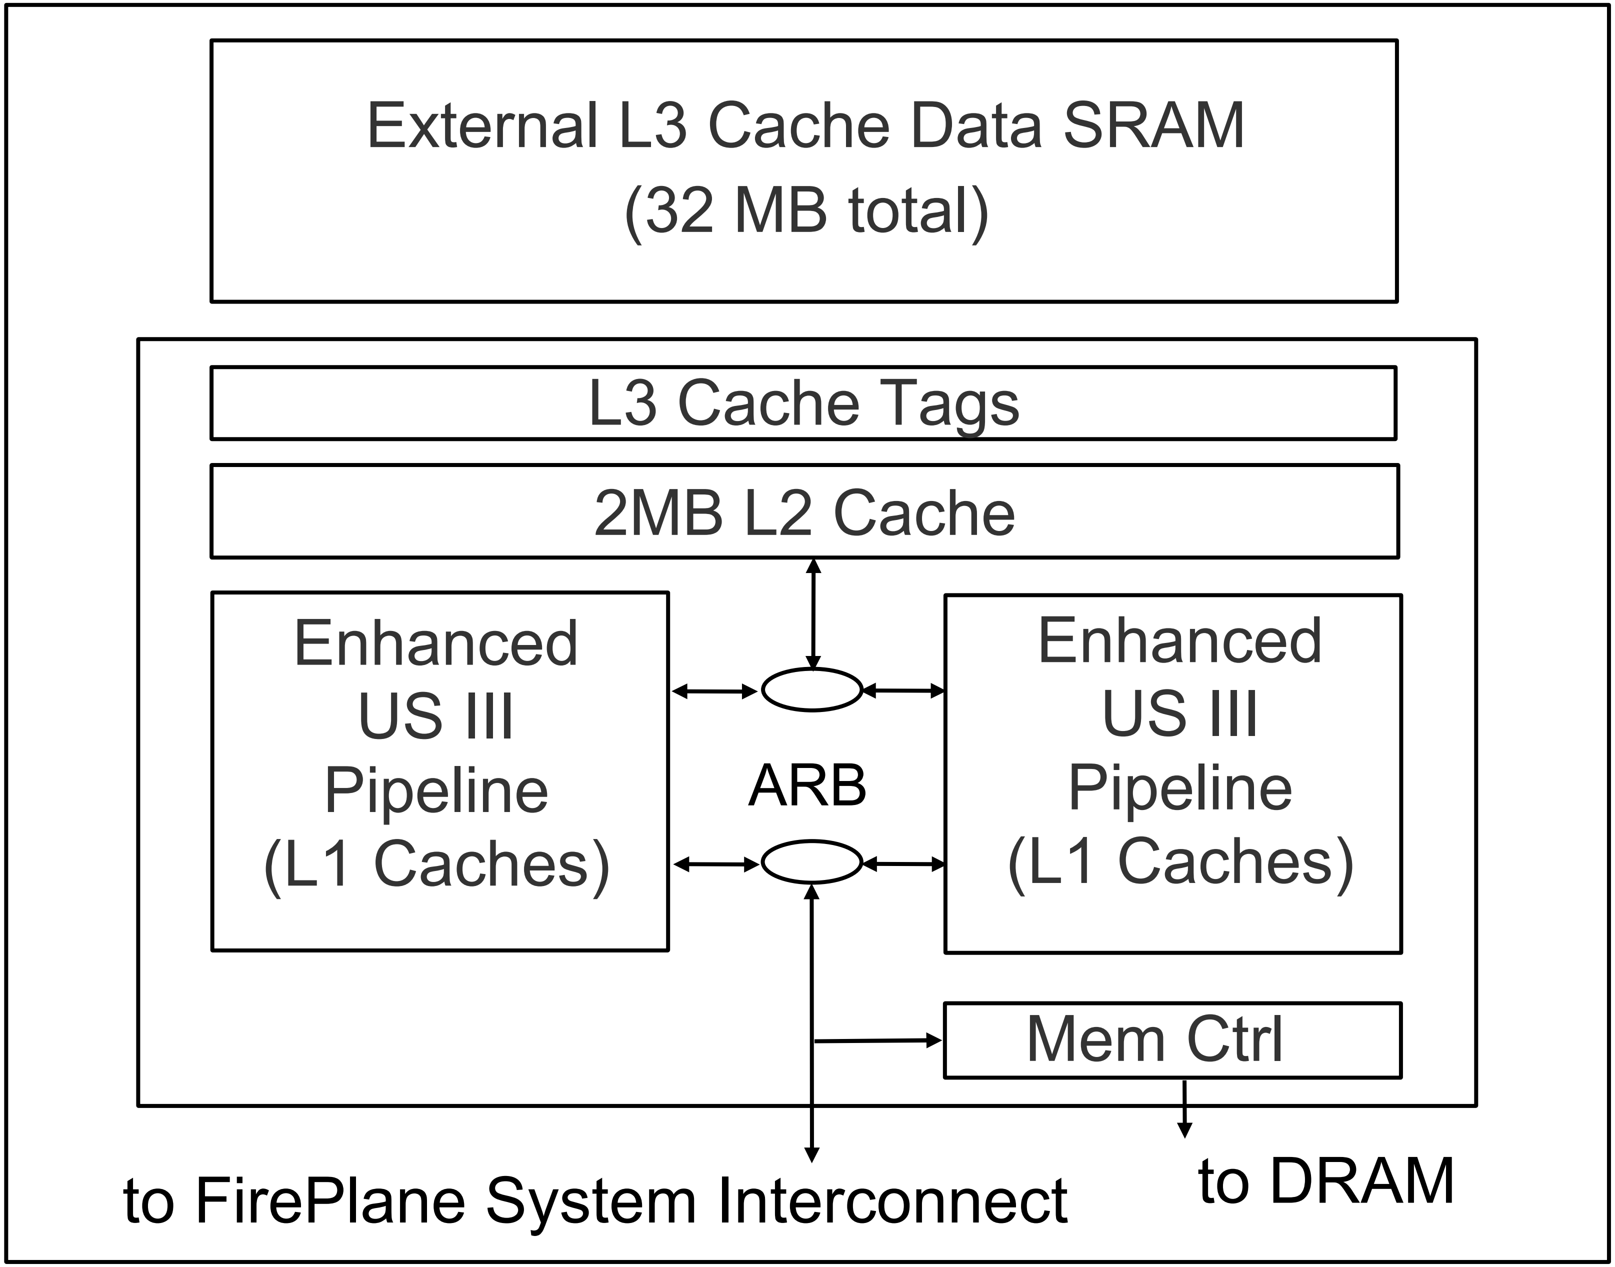

External L3 Cache Data SRAM
(32 MB total)

L3 Cache Tags

2MB L2 Cache

Enhanced
US III
Pipeline
(L1 Caches)

ARB

Enhanced
US III
Pipeline
(L1 Caches)

Mem Ctrl

to FirePlane System Interconnect

to DRAM

# Comparison with Niagara

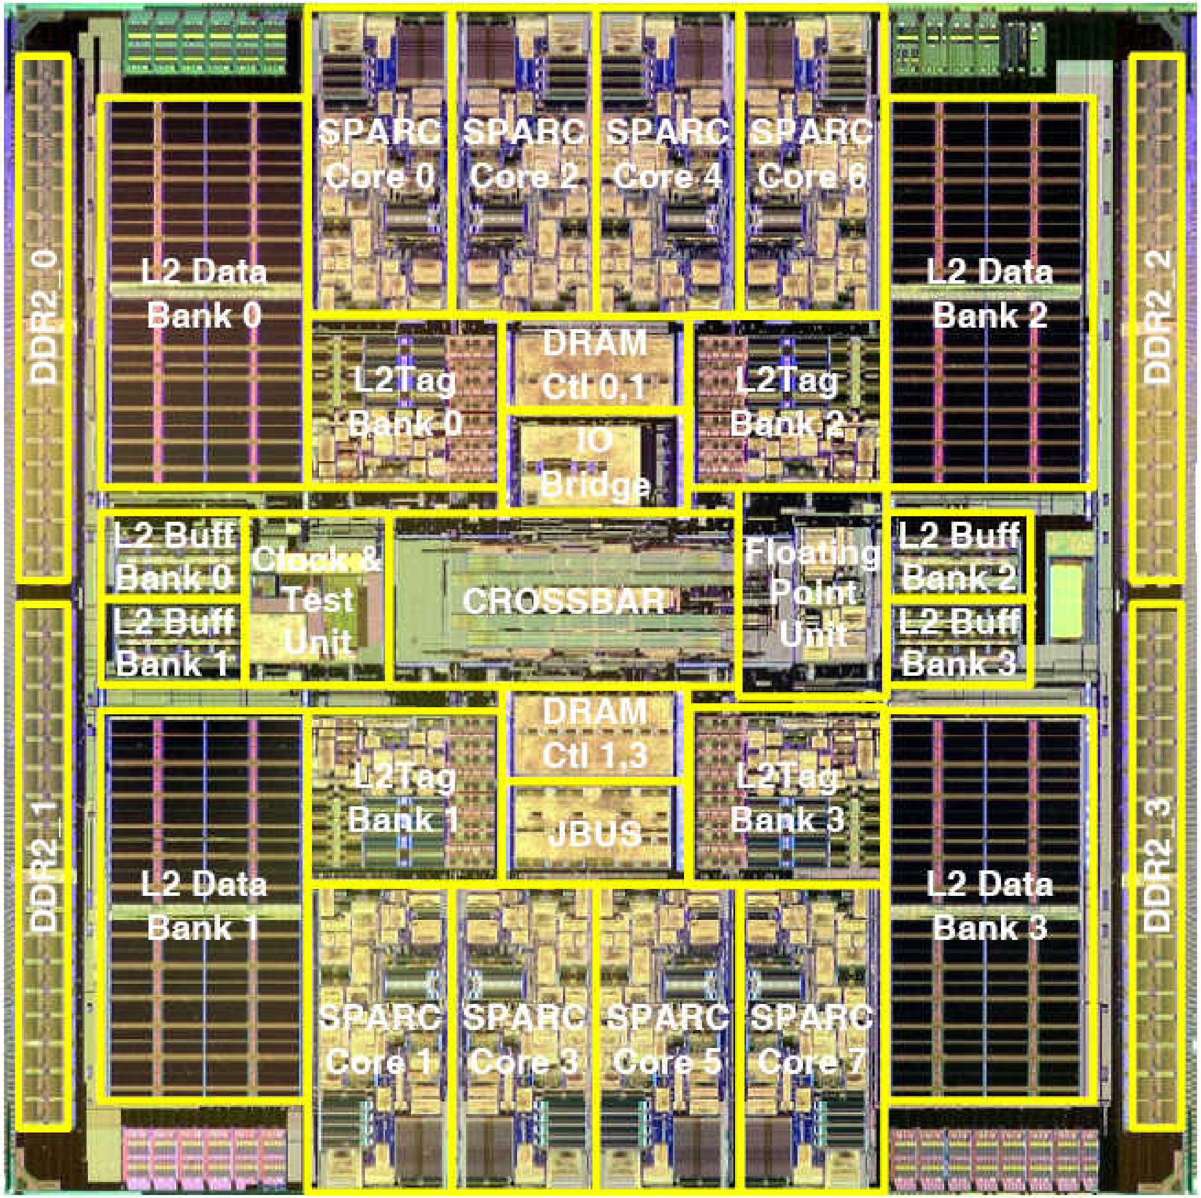

# Physical Implementation

- Use TI's 90nm dual-Vt, dual-gate-oxide technology with 9 layers of low-k dielectric metal
- Increase core frequency beyond shrink factor
- Aggressively manage leakage and power
- Maximize reliability and reduce process defects
- Facilitate rapid block composition
- Improve clocking and reduce skew
- Required 100% recomposition of all blocks

# Library Structure and Usage

- Single data and control library for simplicity
- Library cells extensively used for custom design
- Consistent library characterization
- Built in signal shielding to manage noise
- Local clocks and nwell bias accounted for in template
- Global methodology for substrate and nwell bias control
- Cell optimization to maximize speed and reduce power

# Library Metal Structure

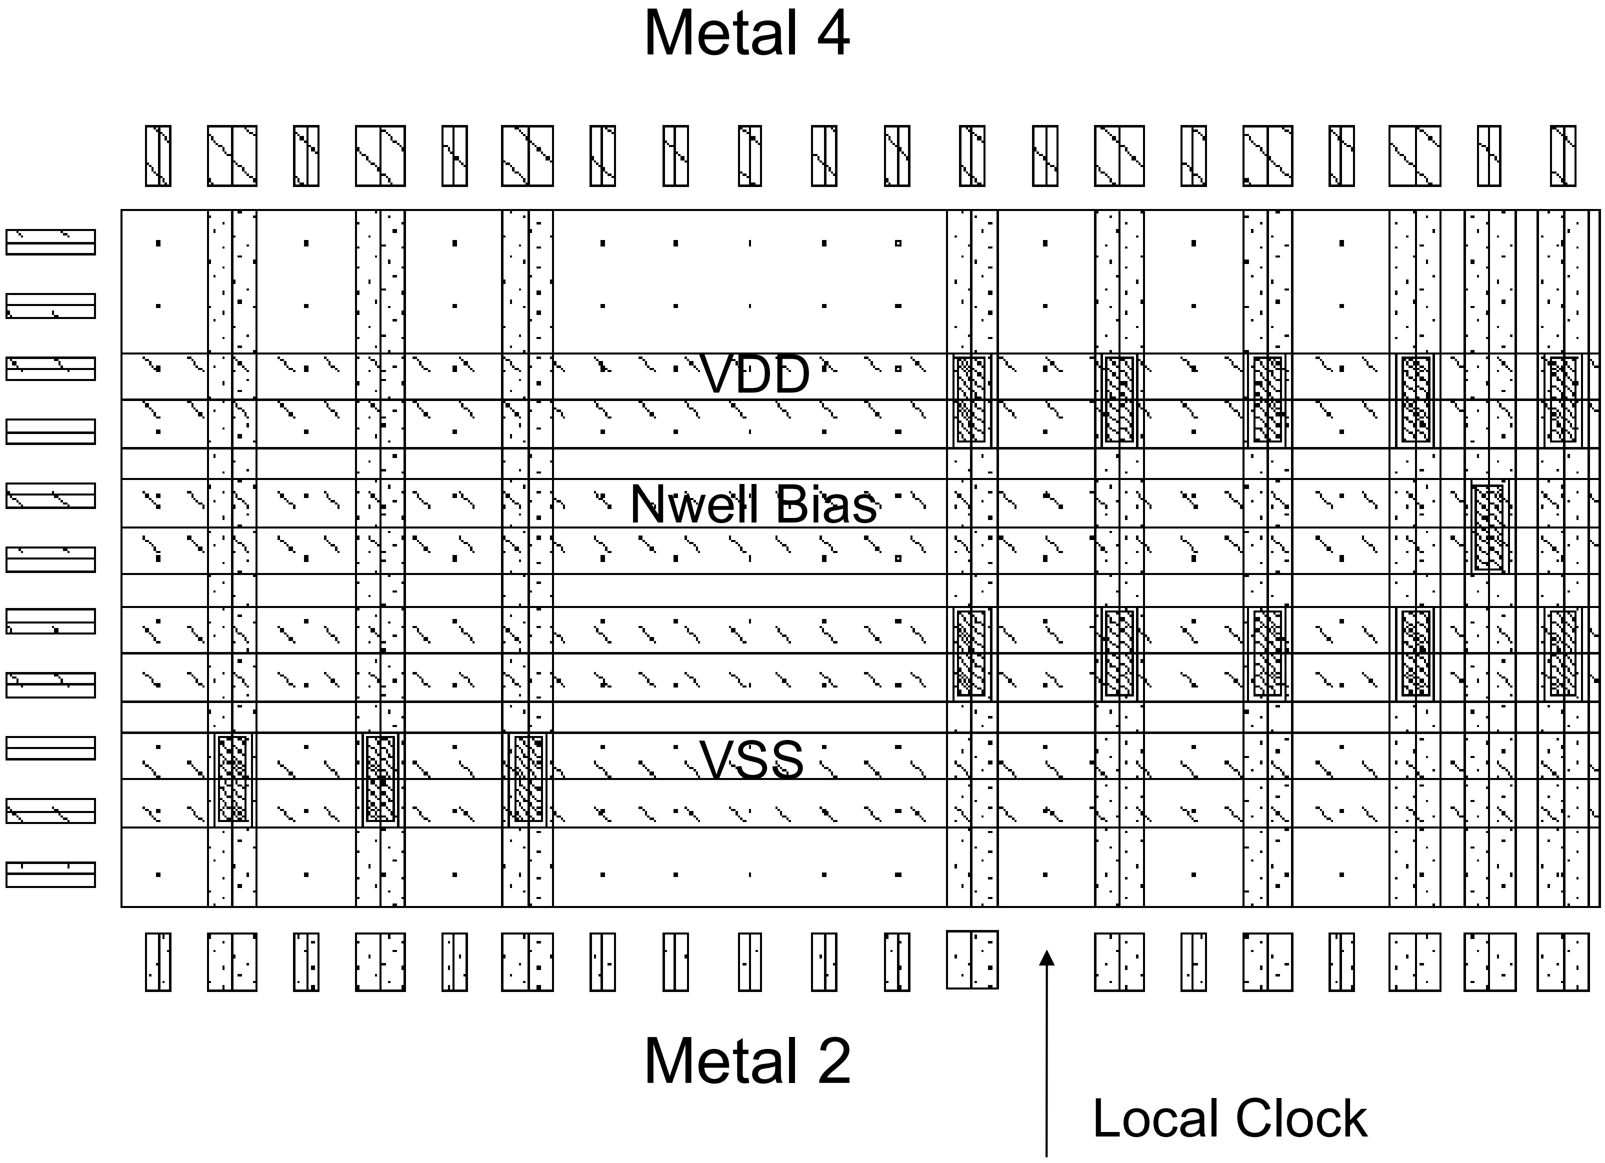

# Power and Leakage Management

- Use gated clocks to dynamically shut off blocks
- Replace high-speed staticized dynamic flip-flops with low power master-slave flip-flops
- Optimize drivers and gates in non-timing critical paths
- Minimize low-Vt transistors to less than 5%
- Modulate body bias to reduce leakage
- Global and local decoupling capacitor insertion resulting in 450nF total chip capacitance

# Leakage vs Vdd with Different Body Bias

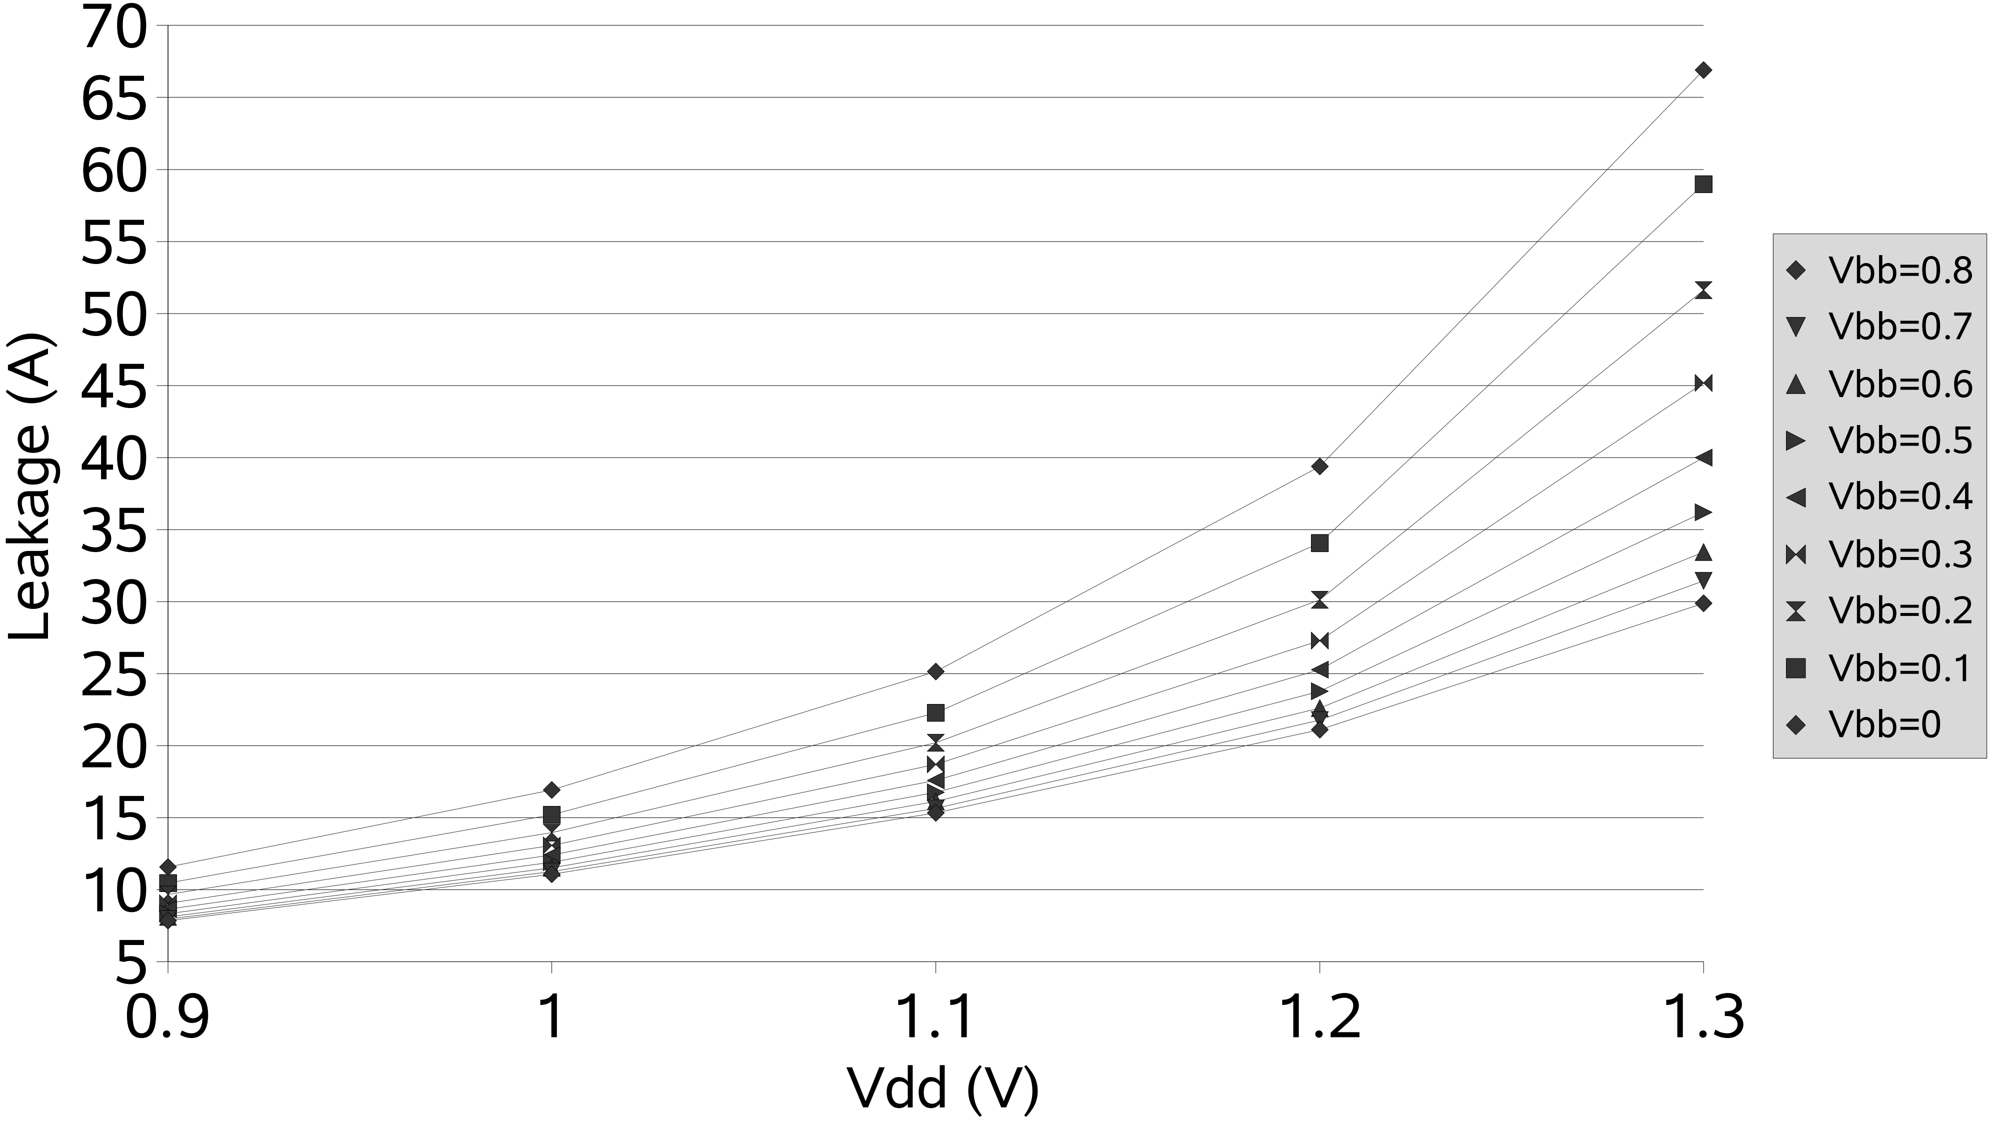

# Flip-flop and Clock Model for Power Down

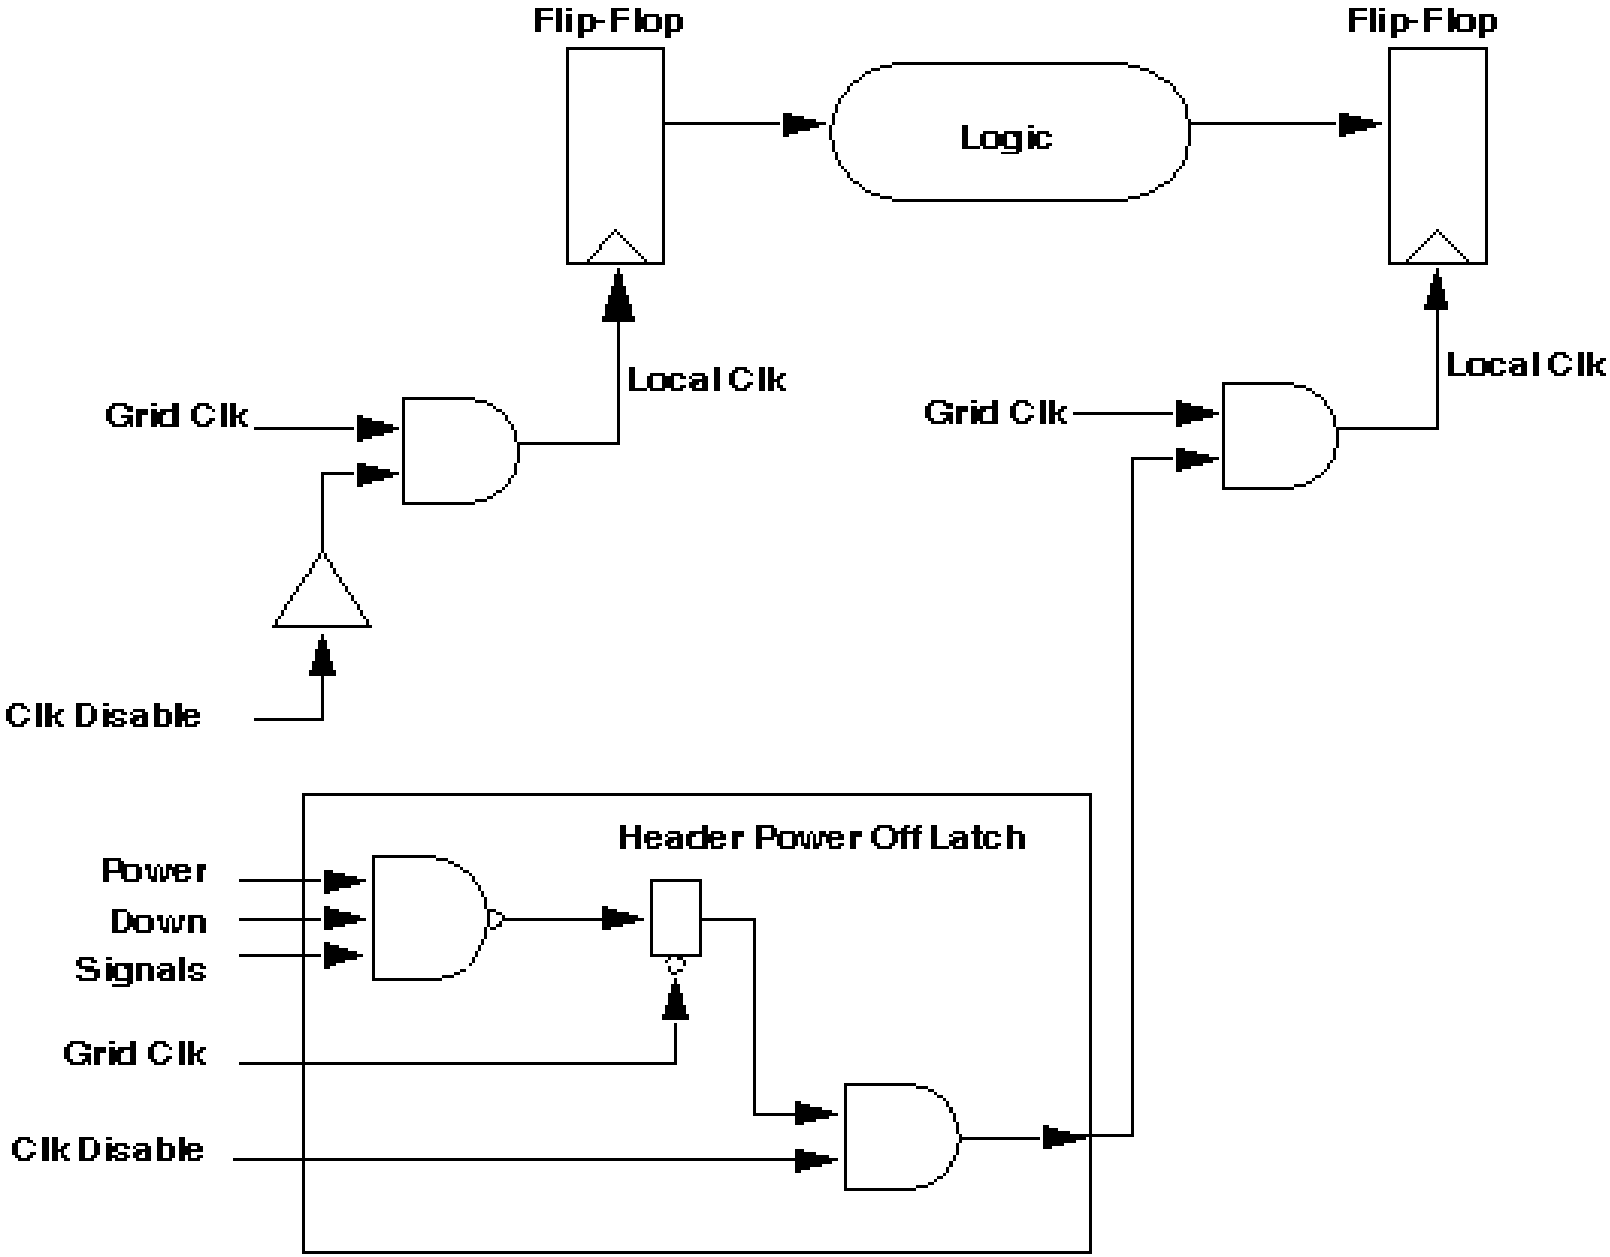

# Chip Power

- 90W @ 1.8GHz @ 1.1v
- >5W reduction using power down feature
- Core (2x): 44%
- EMU: 9%
- L2 + L3: 10%
- Global Clock: 6%
- CPU Route: 18%
- IO: 3%
- Leakage: 10%

**Core Power Profile**

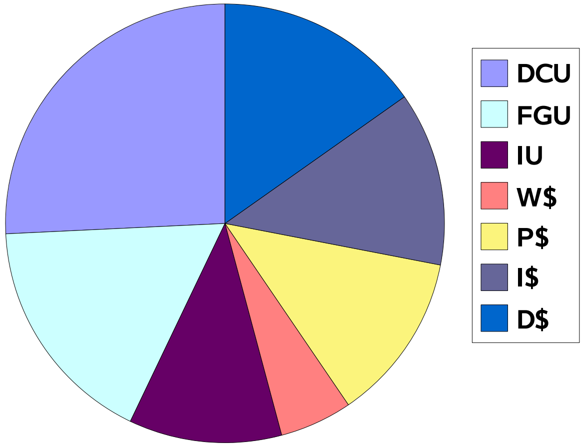

**Fullchip Power Profile**

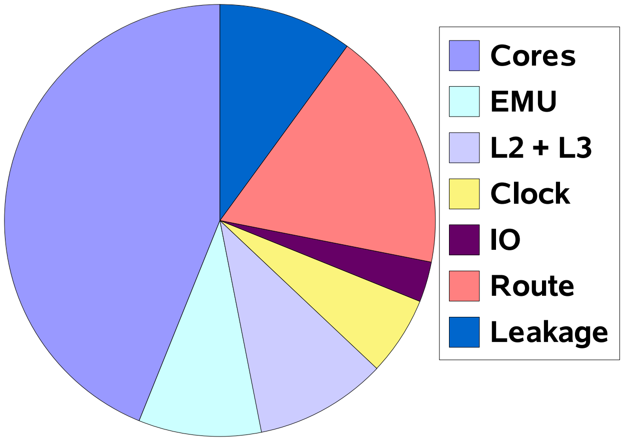

# Fullchip Integration

- Use fully shielded interconnect to eliminate capacitive and inductive noise
- Pre-insert repeaters and decoupling capacitors
- Shields inserted last, allowing better resource utilization for vias and jogs
- Extensive use of area pins for better timing and reduced congestion
- Block level pins matched to integration metal structure

# Core Floorplan

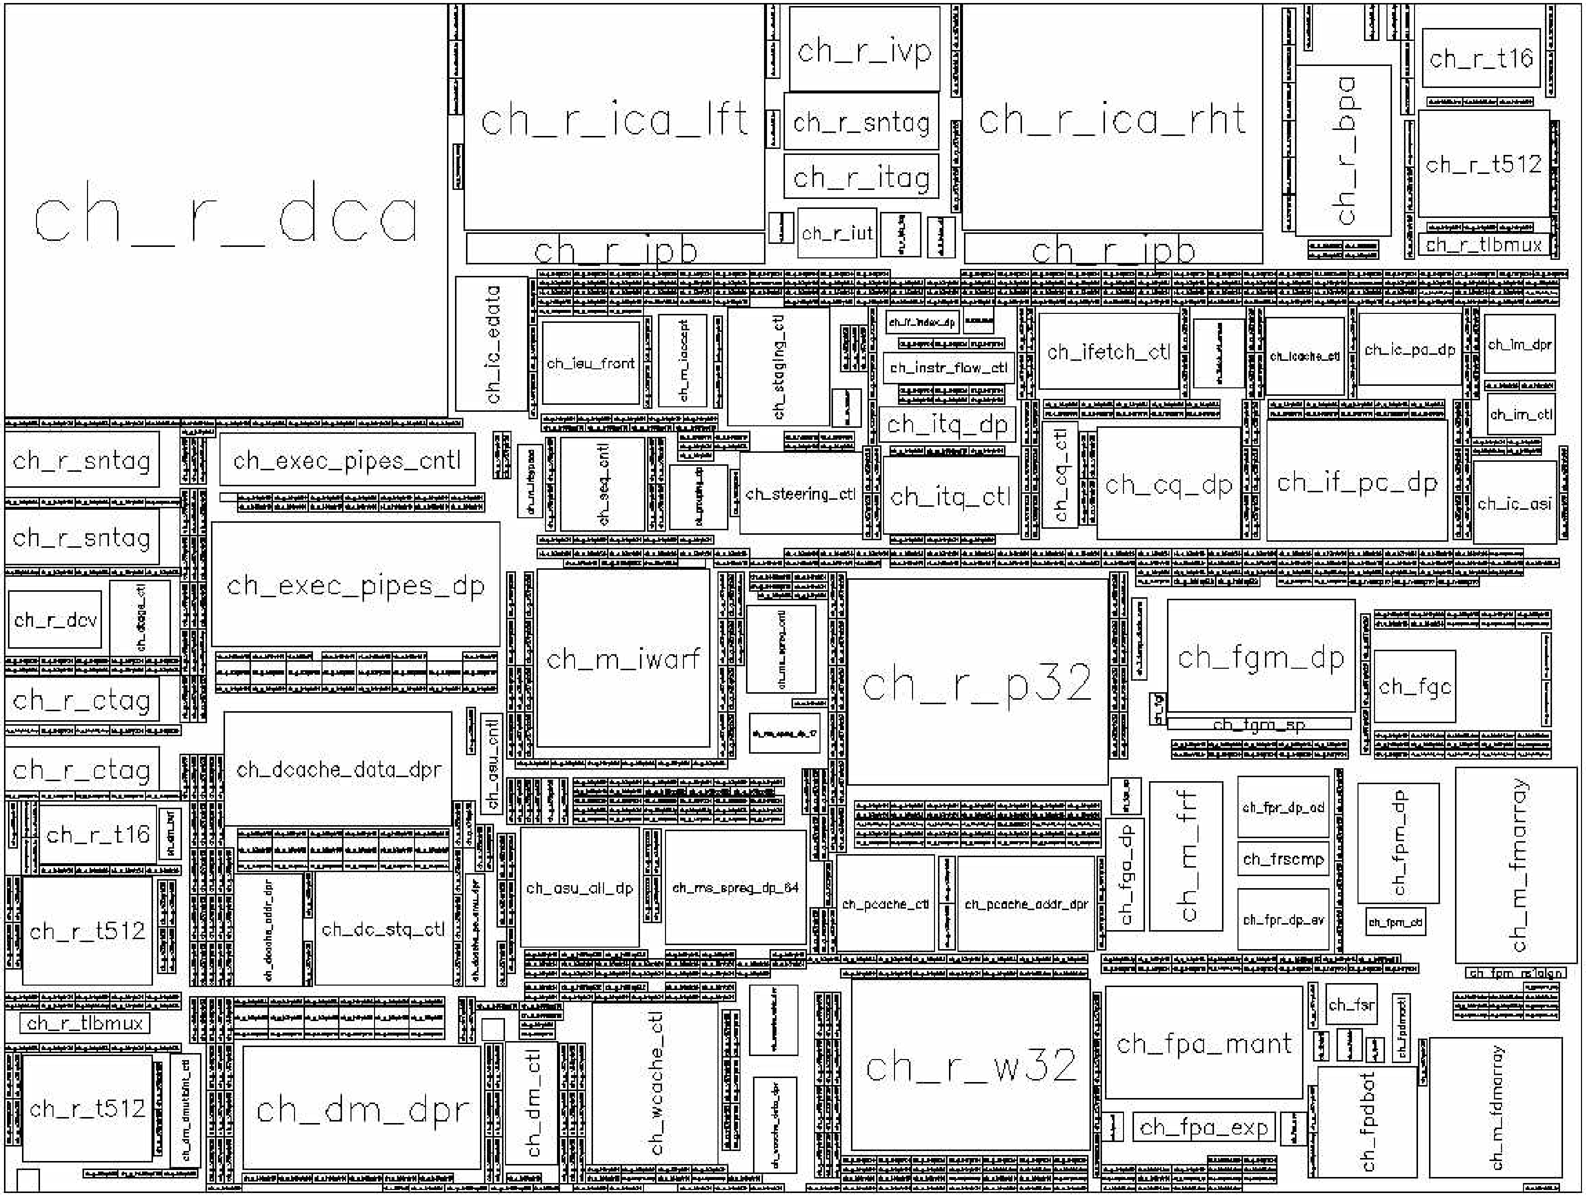

# Clock Generation and Distribution

- 16-stage buffer tree distributes clock from PLL to global grid
- Global M5/M6 grid reduces skew and simplifies clock distribution
- Local block clocks distributed in the data direction from rows of headers
- Built in clock disable for final header control
- PLL allows insertion of full speed clock in scan mode
- PLL mixes ½ speed cycles into normal operation for timing path debug

# Clock Multiplexing for Timing Debug

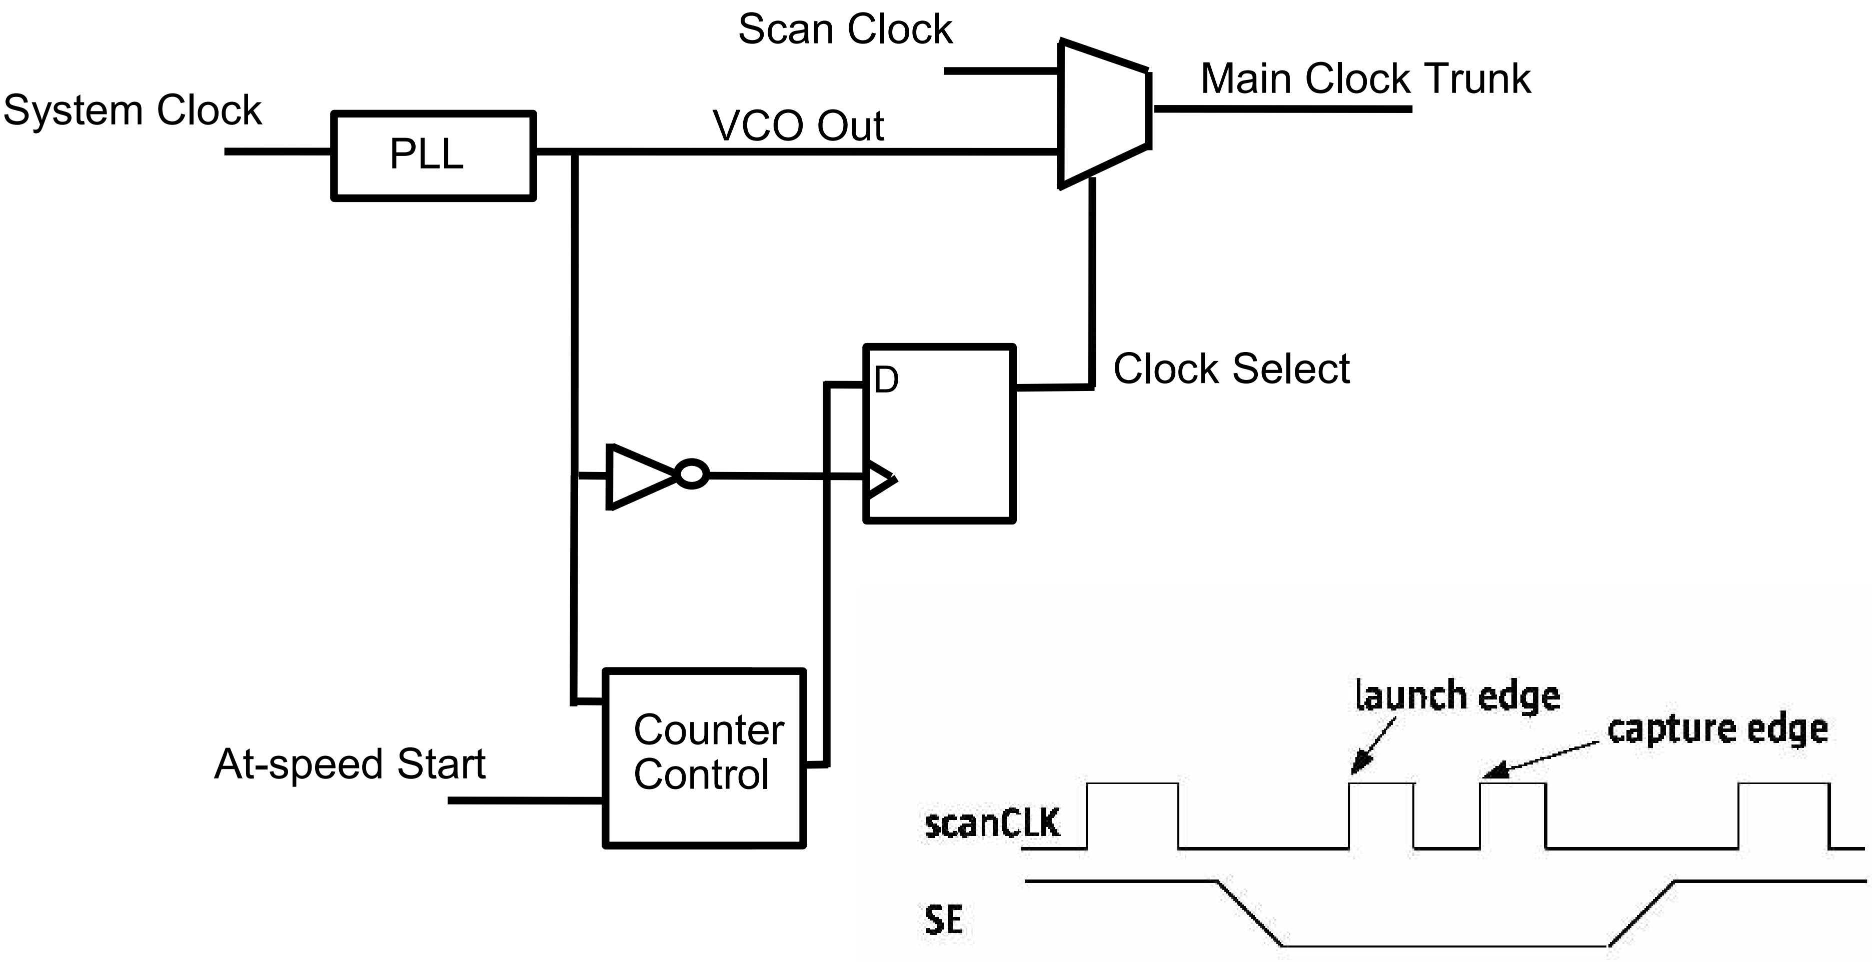

# Rise Delay Distribution at Clock Grid (in ps)

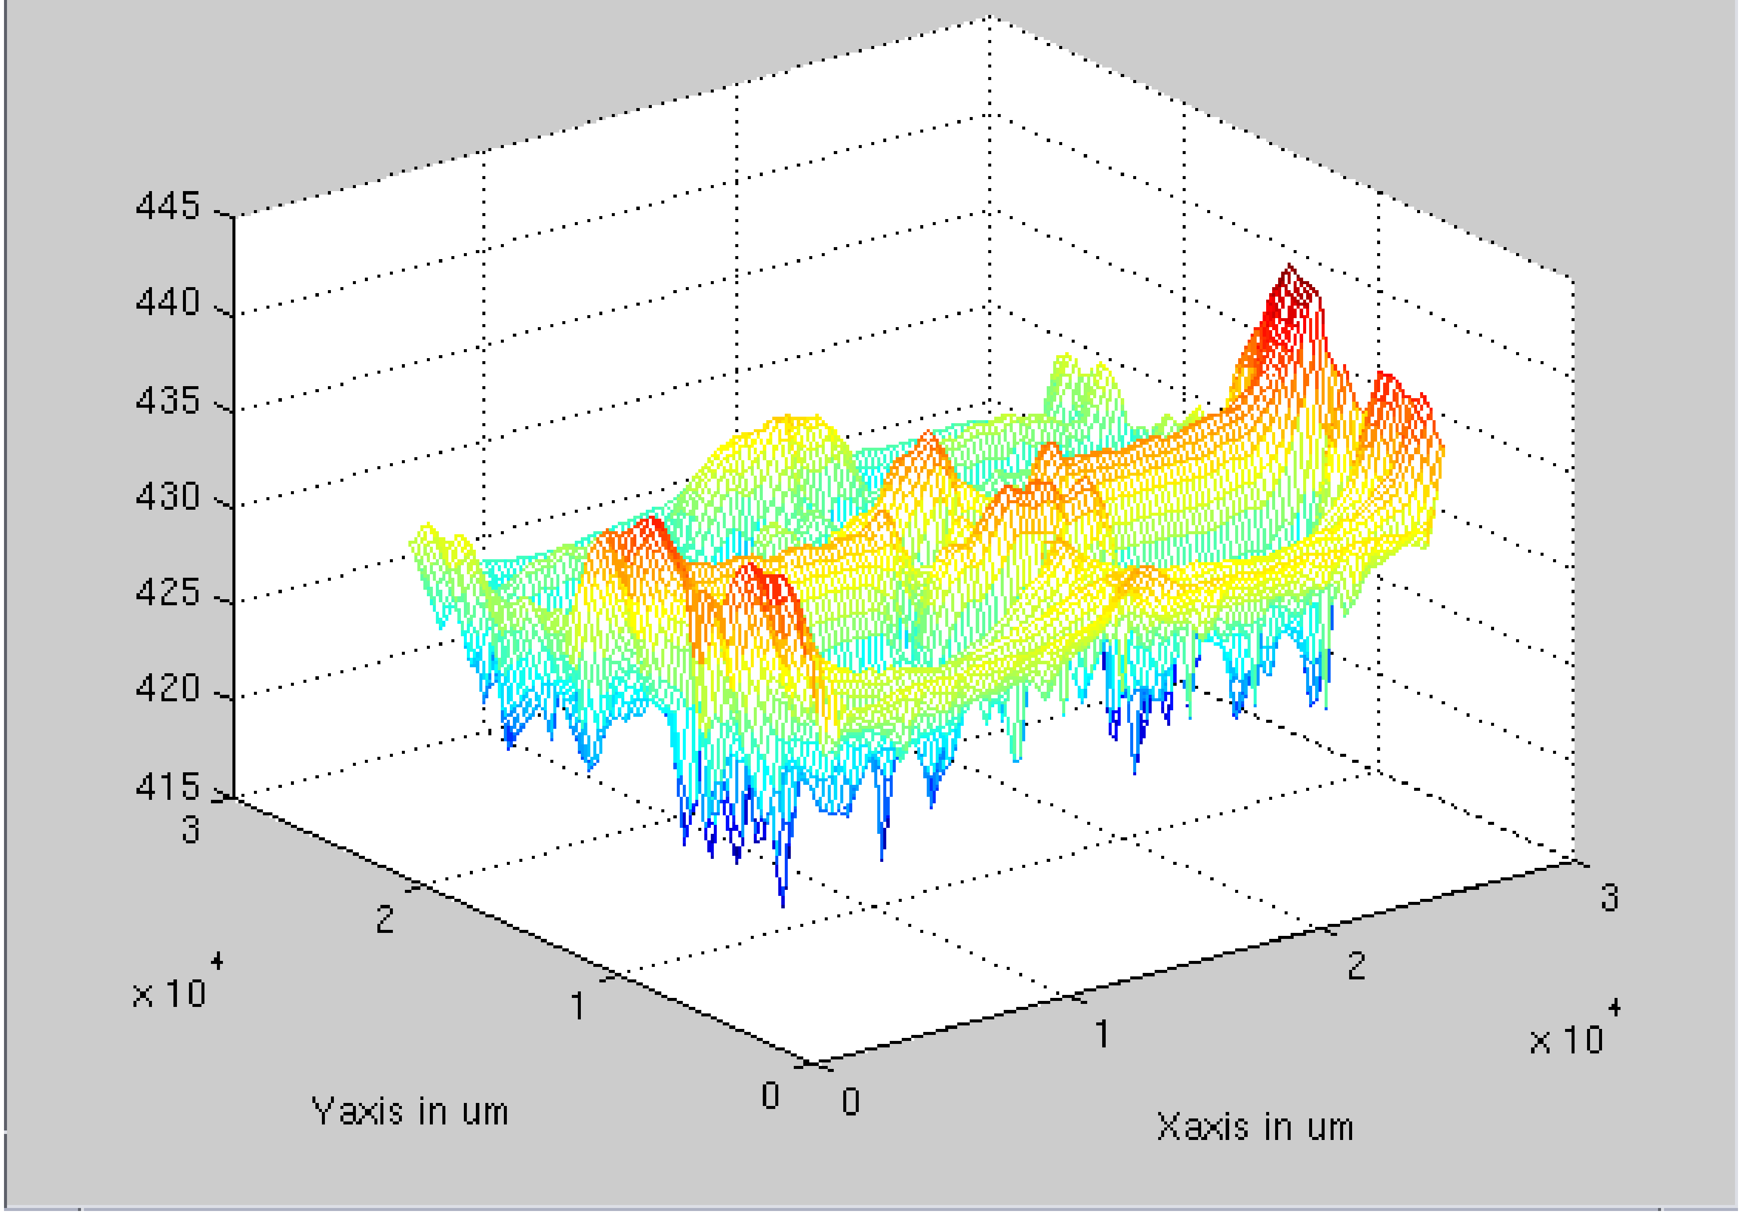

# IO Enhancements

- Extended input common-mode range of the receivers (0.5V to 1.4V)
- Pseudo-differential DTL receiver with shared voltage reference line
- Specified to resolve a minimum of 100mV of voltage differential
- Statistical modeling for the devices and mismatch terms
- Clean power for shielding obtained from off chip

# Memory Design

- Switched from self-timed to frequency-dependent
- First half of cycle for read and write access
- Second half of cycle for bitline sensing and precharge/equilibrate
- Control and address inputs are converted to half-cycle pulses using dynamic flip-flops
- Combine static, dynamic and self-resetting gates
- Self resetting gates used to lock signals together
- Register files use static sensing of single ended bitlines

# Sense-Amp Design

- Common sense-amp design for all SRAM's
- Gate-fed differential sense amplifier replaces drain fed
- Reduced bitline load
- Isolated sense nodes

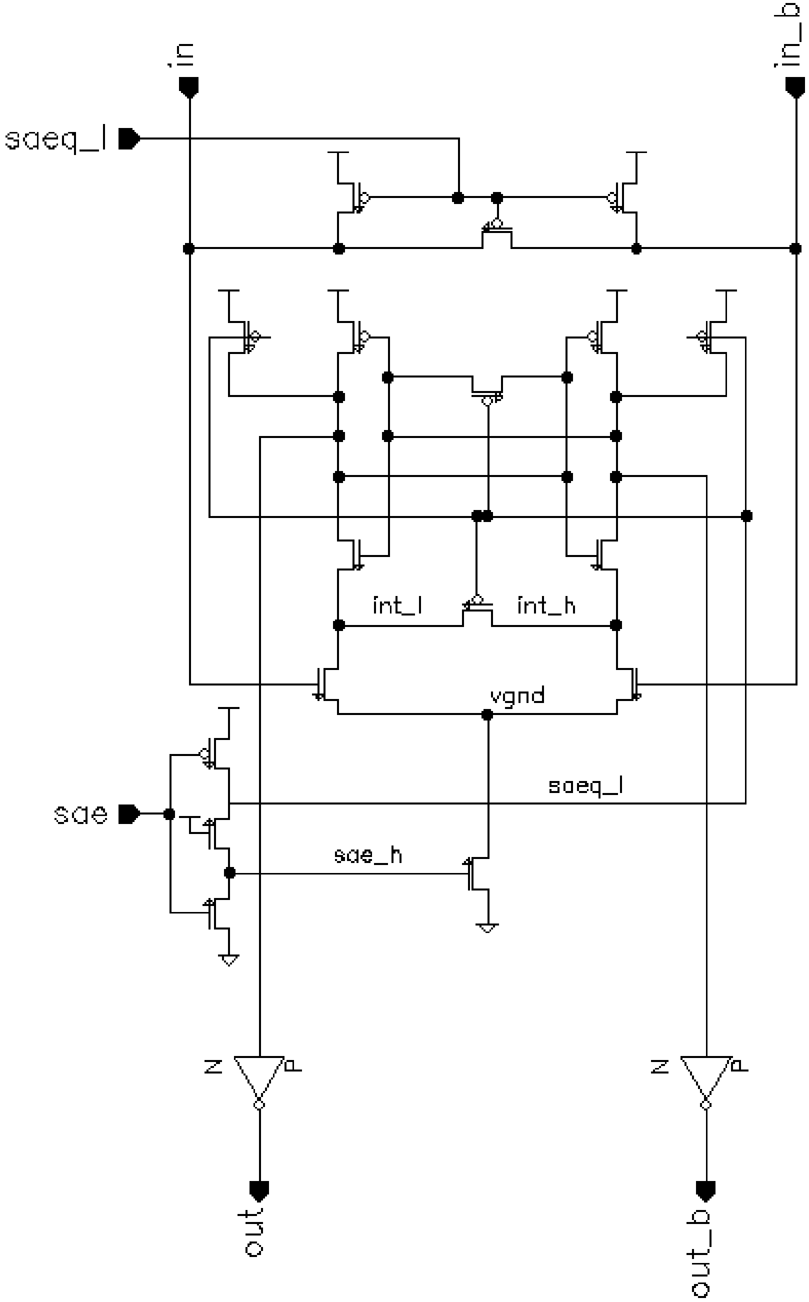

# Preproductization Test Chip Activities

- Layout techniques for better yield
- Strained silicon
- Vmin shift in SRAMs
- Challenges in memory testing
- Aligning performance of various devices

# Layout Techniques

- Significant interaction between layout and what is actually manufactured on silicon
- Meeting design rules is not enough
- Problem with standard mask design practice of packing things up in the smallest possible area
  - Mask designers are trained for this!
- Second pass at layouts
  - Pull geometries away; cleanup to reduce corners
  - Can be done with very little or no area hit
- Gives big improvement in expected yield

# Strained silicon

- Industry wide initial hesitation with defects but universally adopted now
- Idrive gets about half the percentage increase in mobility
- NMOS: tensile (ex. Cap layer, STI, spacer)
  PMOS: compressive (ex. SiGe S/D, metal gate)
- Stress management for silicon
  - Strain and hence device performance will be function of device size and surrounding geometry
  - Statistical variability needs to be better understood
  - Defect density ?

# Vmin shift in SRAMs

Pre-Burnin

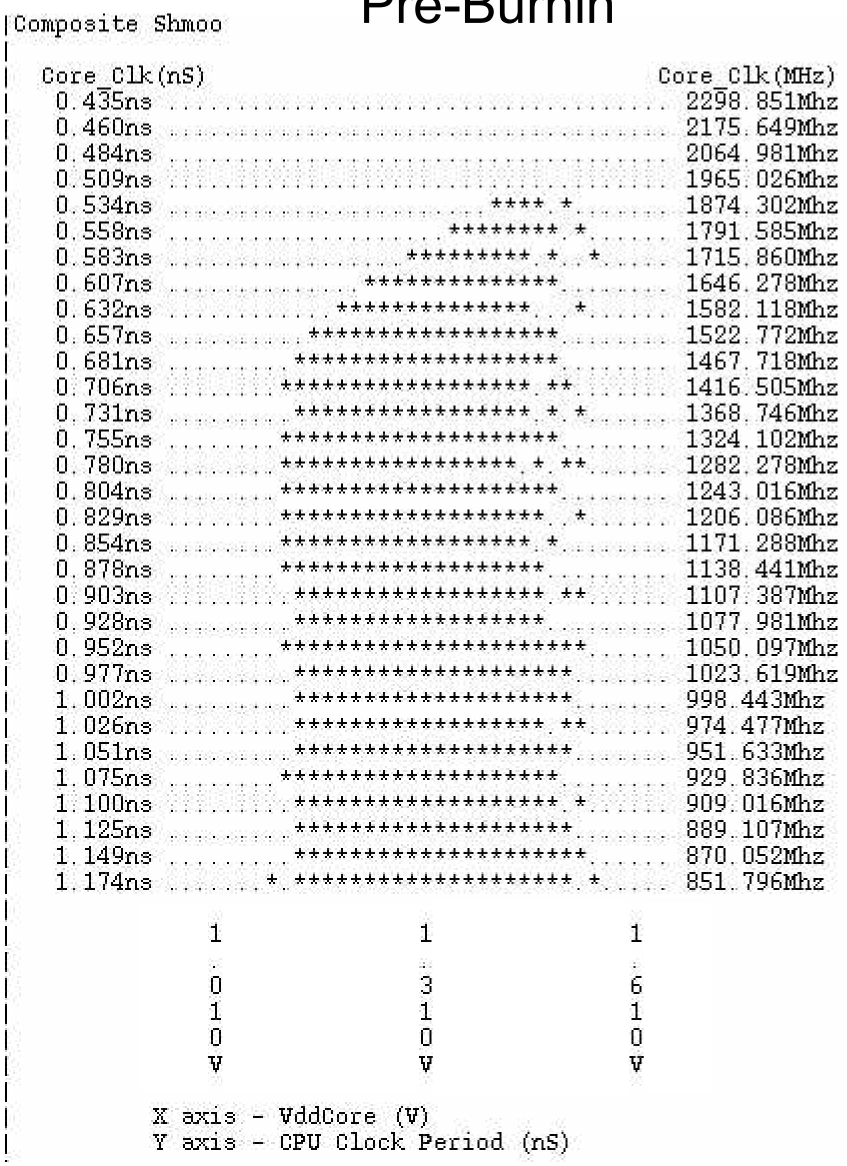

```
|Composite Shmoo
|
|  Core_Clk(nS)                                Core_Clk(MHz)
|   0.435ns .................................... 2298.851Mhz
|   0.460ns .................................... 2175.649Mhz
|   0.484ns .................................... 2064.981Mhz
|   0.509ns .................................... 1965.026Mhz
|   0.534ns .......................****.*....... 1874.302Mhz
|   0.558ns ....................********.*...... 1791.585Mhz
|   0.583ns .................*********.*..*..... 1715.860Mhz
|   0.607ns ..............**************........ 1646.278Mhz
|   0.632ns ............**************...*...... 1582.118Mhz
|   0.657ns ..........******************........ 1522.772Mhz
|   0.681ns .........*******************........ 1467.718Mhz
|   0.706ns ........******************.**....... 1416.505Mhz
|   0.731ns .........*****************.*.*...... 1368.746Mhz
|   0.755ns ........********************........ 1324.102Mhz
|   0.780ns ........*****************.*.**...... 1282.278Mhz
|   0.804ns ........********************........ 1243.016Mhz
|   0.829ns ........*******************..*...... 1206.086Mhz
|   0.854ns ........******************.*........ 1171.288Mhz
|   0.878ns ........*******************......... 1138.441Mhz
|   0.903ns .........******************.**...... 1107.387Mhz
|   0.928ns .........******************......... 1077.981Mhz
|   0.952ns ........**********************...... 1050.097Mhz
|   0.977ns .........********************....... 1023.619Mhz
|   1.002ns .........********************....... 998.443Mhz
|   1.026ns .........******************.**...... 974.477Mhz
|   1.051ns .........********************....... 951.633Mhz
|   1.075ns ........********************........ 929.836Mhz
|   1.100ns .........*******************.*...... 909.016Mhz
|   1.125ns .........********************....... 889.107Mhz
|   1.149ns .........*********************...... 870.052Mhz
|   1.174ns .......*.********************.*..... 851.796Mhz
|
|              1              1              1
|              .              .              .
|              0              3              6
|              1              1              1
|              0              0              0
|              V              V              V
|
|          X axis - VddCore (V)
|          Y axis - CPU Clock Period (nS)
```

Post-Burnin

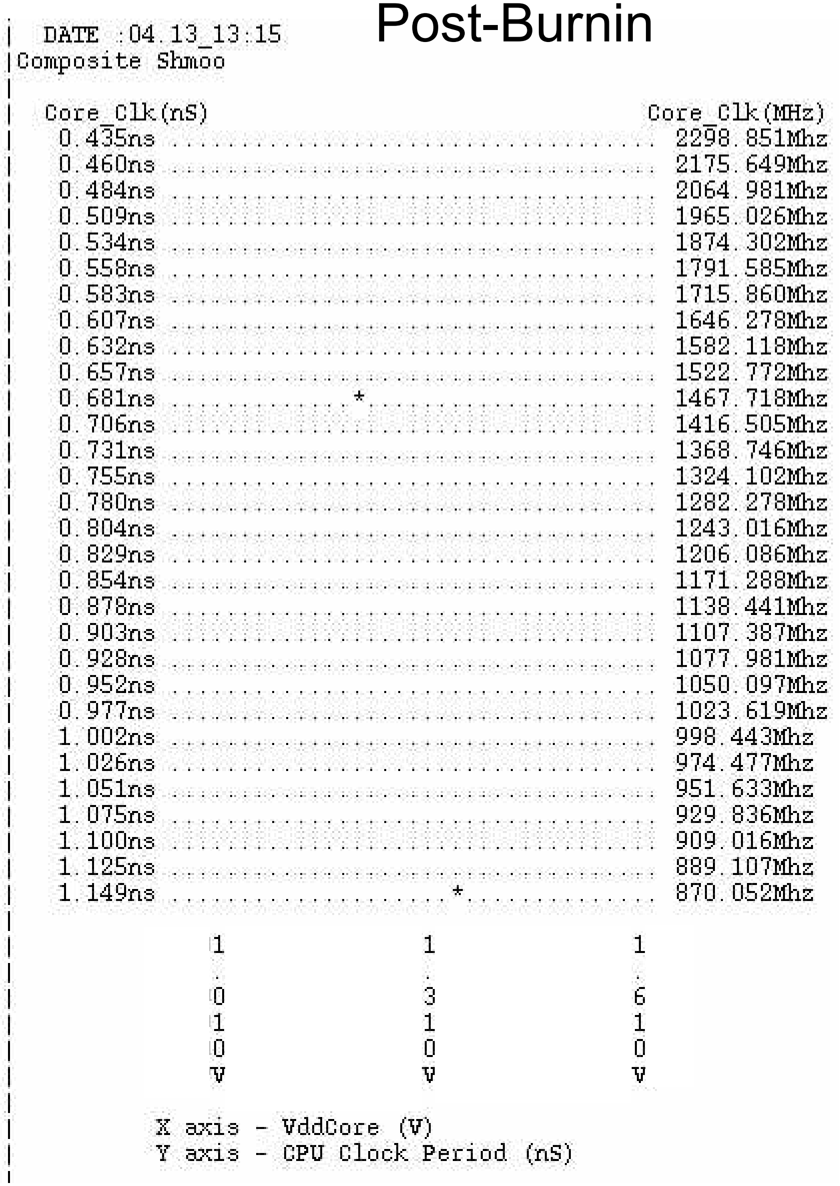

```
|  DATE  :04.13_13:15
|Composite Shmoo
|
|  Core_Clk(nS)                                Core_Clk(MHz)
|   0.435ns .................................... 2298.851Mhz
|   0.460ns .................................... 2175.649Mhz
|   0.484ns .................................... 2064.981Mhz
|   0.509ns .................................... 1965.026Mhz
|   0.534ns .................................... 1874.302Mhz
|   0.558ns .................................... 1791.585Mhz
|   0.583ns .................................... 1715.860Mhz
|   0.607ns .................................... 1646.278Mhz
|   0.632ns .................................... 1582.118Mhz
|   0.657ns .................................... 1522.772Mhz
|   0.681ns .............*...................... 1467.718Mhz
|   0.706ns .................................... 1416.505Mhz
|   0.731ns .................................... 1368.746Mhz
|   0.755ns .................................... 1324.102Mhz
|   0.780ns .................................... 1282.278Mhz
|   0.804ns .................................... 1243.016Mhz
|   0.829ns .................................... 1206.086Mhz
|   0.854ns .................................... 1171.288Mhz
|   0.878ns .................................... 1138.441Mhz
|   0.903ns .................................... 1107.387Mhz
|   0.928ns .................................... 1077.981Mhz
|   0.952ns .................................... 1050.097Mhz
|   0.977ns .................................... 1023.619Mhz
|   1.002ns .................................... 998.443Mhz
|   1.026ns .................................... 974.477Mhz
|   1.051ns .................................... 951.633Mhz
|   1.075ns .................................... 929.836Mhz
|   1.100ns .................................... 909.016Mhz
|   1.125ns .................................... 889.107Mhz
|   1.149ns ....................*............... 870.052Mhz
|
|              1              1              1
|              .              .              .
|              0              3              6
|              1              1              1
|              0              0              0
|              V              V              V
|
|          X axis - VddCore (V)
|          Y axis - CPU Clock Period (nS)
```

# Vmin shift in SRAMs

- Picoprobe data shows up to 150mV Vt mismatch in the driver devices while the standard deviation is 19mV. Adjacent devices differ by up to 8 sigma!!
- Additional bit fails explained by the NBTI shift of ~60-80mV
- Containment by better wafer sort/bit repair methodology
- More details in VLSI Symposium 2006

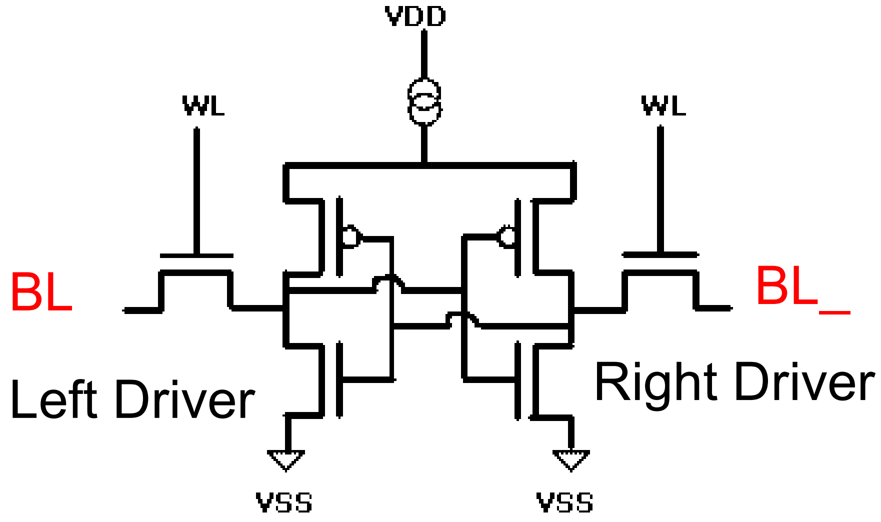

# Challenges in memory test

- Logically ordered embedded memory testing no longer enough
  - Bits physically scrambled due to design and soft error issues
- Test order needs to take physical bit location into account
- New tests revealed failing bits physically adjacent to repaired bits
  - These were not stuck-at faults but marginal failures dependent on the sequence and speed in which bits around the physical defect were addressed
  - Important to have this capability as part of the DFT requirement

# Aligning performance of devices

- Transistors in different flavors
  - Standard Vt, Low Vt, Thick oxide devices (and High Vt devices for future nodes)
- Important to understand how these devices would track as the process shifts
- Thick oxide devices left near the slow limit of the specification to be conservative
  - Caused high frequency cutoff due to the slow paths inside sense amplifier in input buffers

# Performance and Statistics

| | UltraSPARC III | UltraSPARCIV | UltraSPARCIV+ |
|---|---|---|---|
| Supply Voltage | 1.5V | 1.35V | 1.1V |
| Frequency | 0.8GHz | 1.2GHz | 1.8GHz |
| Power | 60W | 102W | 90W |
| Die Size | 233mm$^2$ | 352mm$^2$ | 336mm$^2$ |
| Transistor Count | 23M | 66M | 295M |
| Data Cache | 64KB | 64KB (*2) | 64KB (*2) |
| Instruction Cache | 32KB | 32KB (*2) | 64KB (*2) |
| L2 Cache | 8MB (off chip) | 16MB (off chip) | 2MB (on chip) |
| L3 Cache | NO | NO | 32MB (off chip) |
| L3 Tag | NO | NO | YES (on chip) |
| Technology | 150nm | 130nm | 90nm |

# Schmoo Plot of Processor

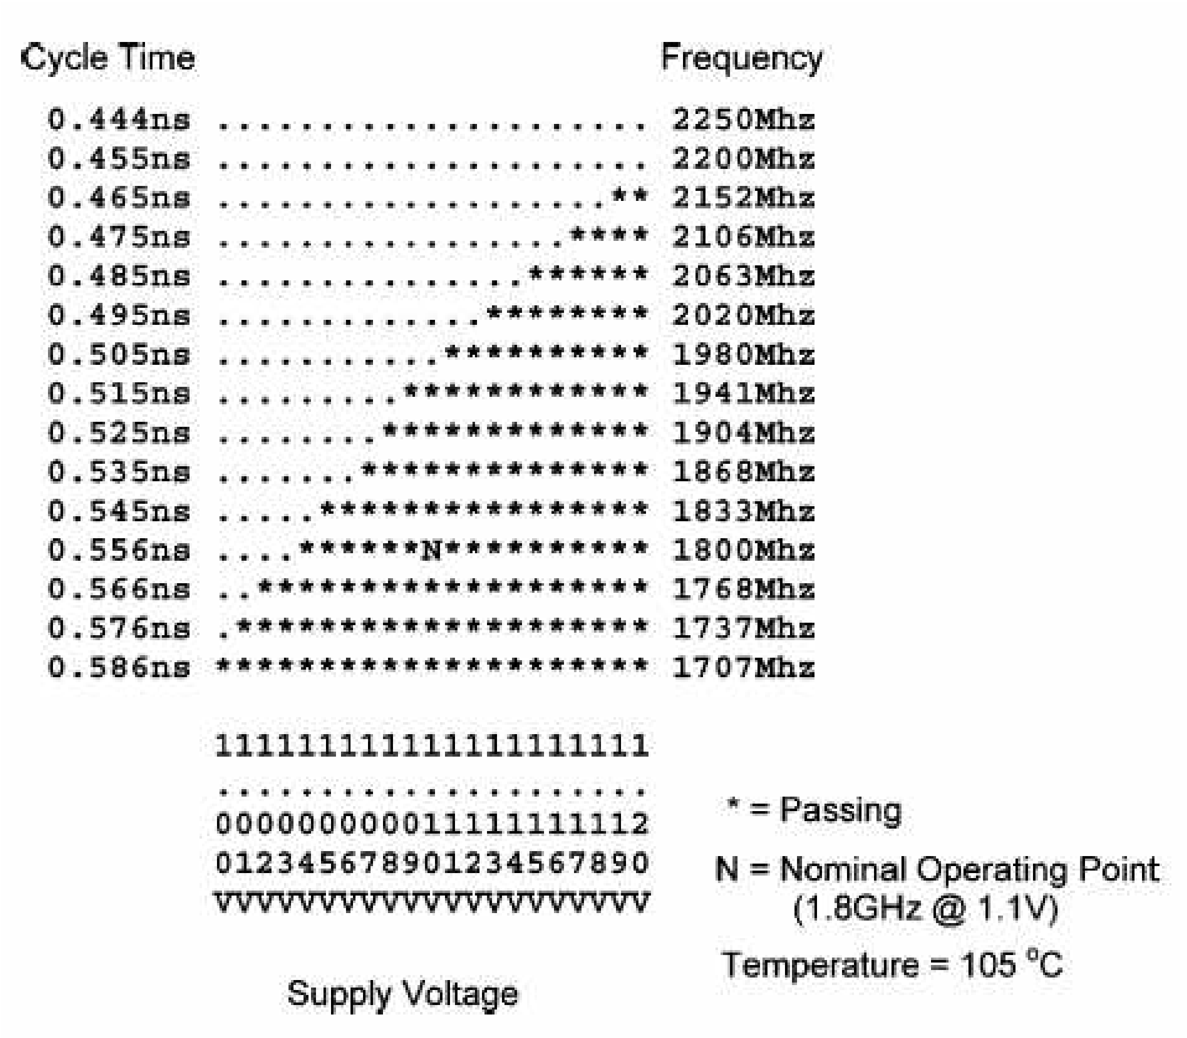

```
Cycle Time                              Frequency

 0.444ns ..................... 2250Mhz
 0.455ns ..................... 2200Mhz
 0.465ns ...................** 2152Mhz
 0.475ns .................**** 2106Mhz
 0.485ns ...............****** 2063Mhz
 0.495ns .............******** 2020Mhz
 0.505ns ...........********** 1980Mhz
 0.515ns .........************ 1941Mhz
 0.525ns ........************* 1904Mhz
 0.535ns .......************** 1868Mhz
 0.545ns .....**************** 1833Mhz
 0.556ns ....******N********** 1800Mhz
 0.566ns ..******************* 1768Mhz
 0.576ns .******************** 1737Mhz
 0.586ns ********************* 1707Mhz

         111111111111111111111
         .....................
         000000000011111111112
         012345678901234567890
         VVVVVVVVVVVVVVVVVVVVV

             Supply Voltage
```

* = Passing

N = Nominal Operating Point
(1.8GHz @ 1.1V)

Temperature = 105 °C

# Micrograph of Processor Die

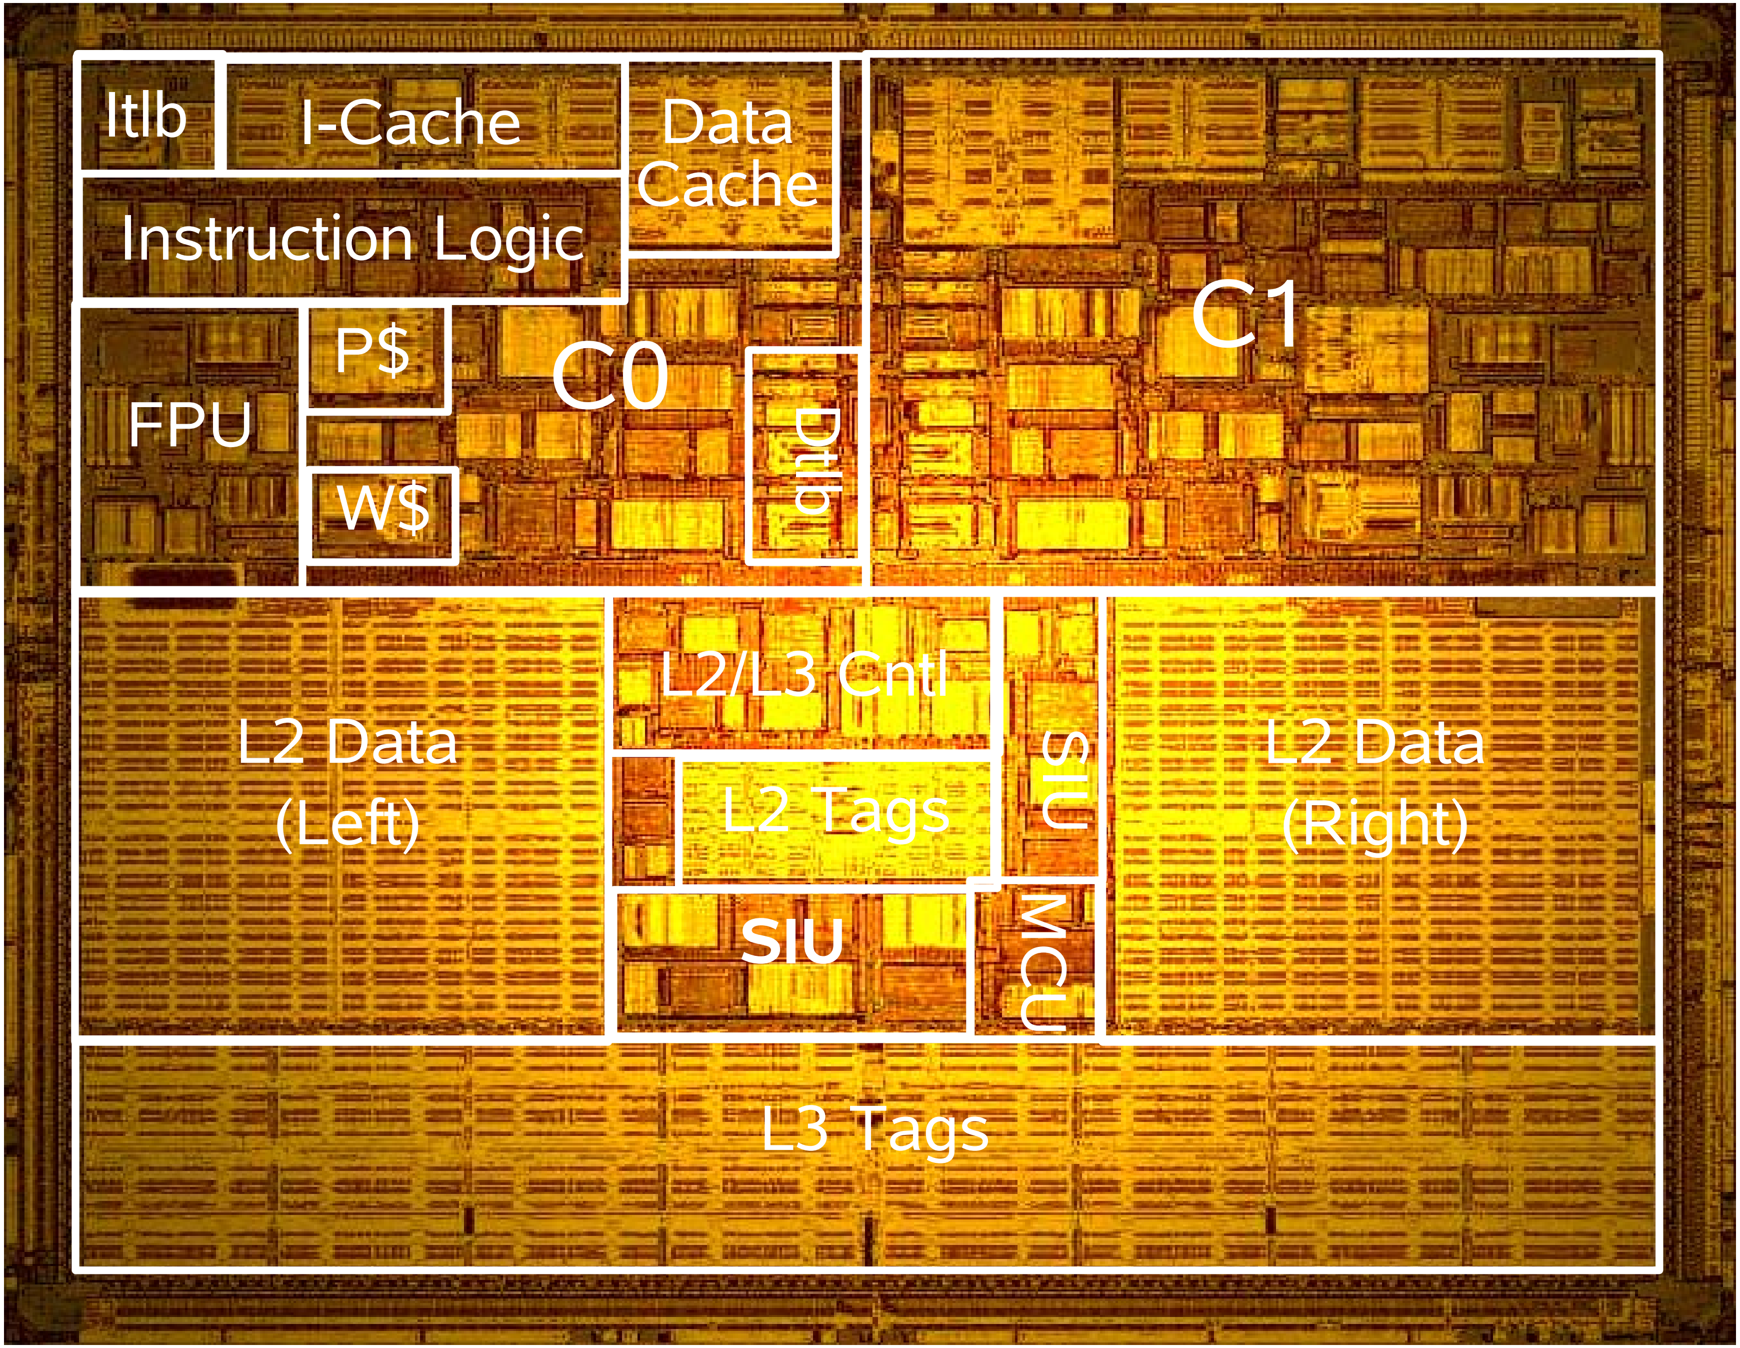

# Implementation of a 4th Generation 1.8GHz Dual-Core SPARC V9 Microprocessor

Anand Dixit, Jason Hart, Swee Yew Choe,
Lik Cheng, Chipai Chou, Kenneth Ho,
Jesse Hsu, Kyung Lee, John Wu

Sun Microsystems, Sunnyvale, CA, USA