

# An Integrated Quad-Core AMD Opteron™ MicroProcessor



S. Searles<sup>1</sup>, J. Dorsey<sup>1</sup>, M. Ciraula<sup>2</sup>, S. Johnson<sup>1</sup>,  
N. Bujanos<sup>1</sup>, D. Wu<sup>1</sup>, M. Braganza<sup>1</sup>, S. Meyers<sup>1</sup>,  
E. Fang<sup>3</sup>, R. Kumar<sup>3</sup>

<sup>1</sup>AMD, Austin, TX; <sup>2</sup>AMD, Fort Collins, CO; <sup>3</sup>AMD, Sunnyvale, CA

# Outline


- Technology Highlights
- Major Features
- Clock/Voltage Domains and PLLs
- Cache Design Choices
- DDR DRAM I/O
- HyperTransport™ (HT) I/O
- Thermal Monitoring
- Mixed Signal Design Study (DDR)

# Technology

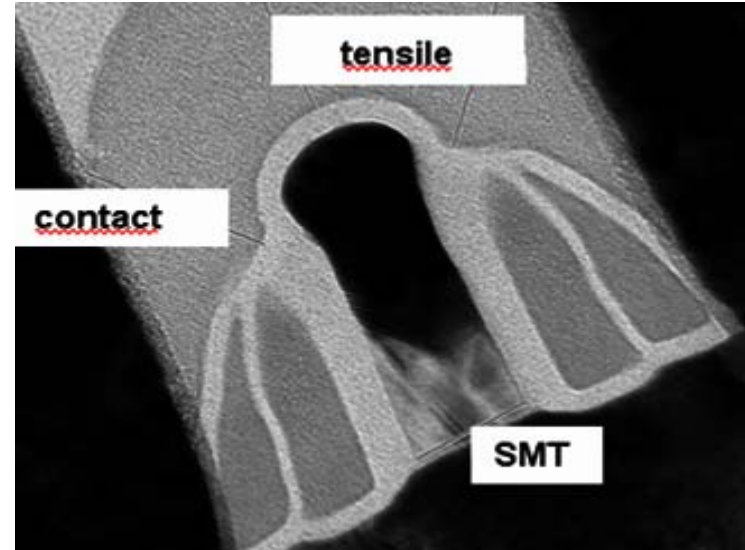
- 65nm, SOI CMOS, 0.8-1.4V
- Dual strain liners + eSiGe
- Floating and Body tied
- 11 Cu layers + low-k
- 1X, 1.3X, 2X, 4X thick Cu
- F04 inv delay, 15ps

Layer	Pitch
Contacted Poly	270nm
M1	200nm
M2	200nm
M3	270nm
M4	270nm
M5	270nm
M6	400nm
M7	400nm
M8	400nm
M9	400nm
M10	800nm
M11	800nm


# Cross Sections



NMOS



PMOS




# Notable Attributes

- Monolithic, Quad x86 cores + Northbridge
- Symmetric cores, same DRAM/HT latency
- Enhanced 128b FPU
- Shared 32Way, 2MB L3
- Each core has an exclusive 16Way, 0.5MB L2
- Cores have 2Way, 64kB L1 I and D Caches
- 2 DDR channels, 4 DIMM/Channel, DDR2+3
- Enhanced HT links for 2 to 8Way MP



# Die Overview



- 463 Million Transistors
- 283mm<sup>2</sup>
- 2.0GHz & up
- 0.8 to 1.4V


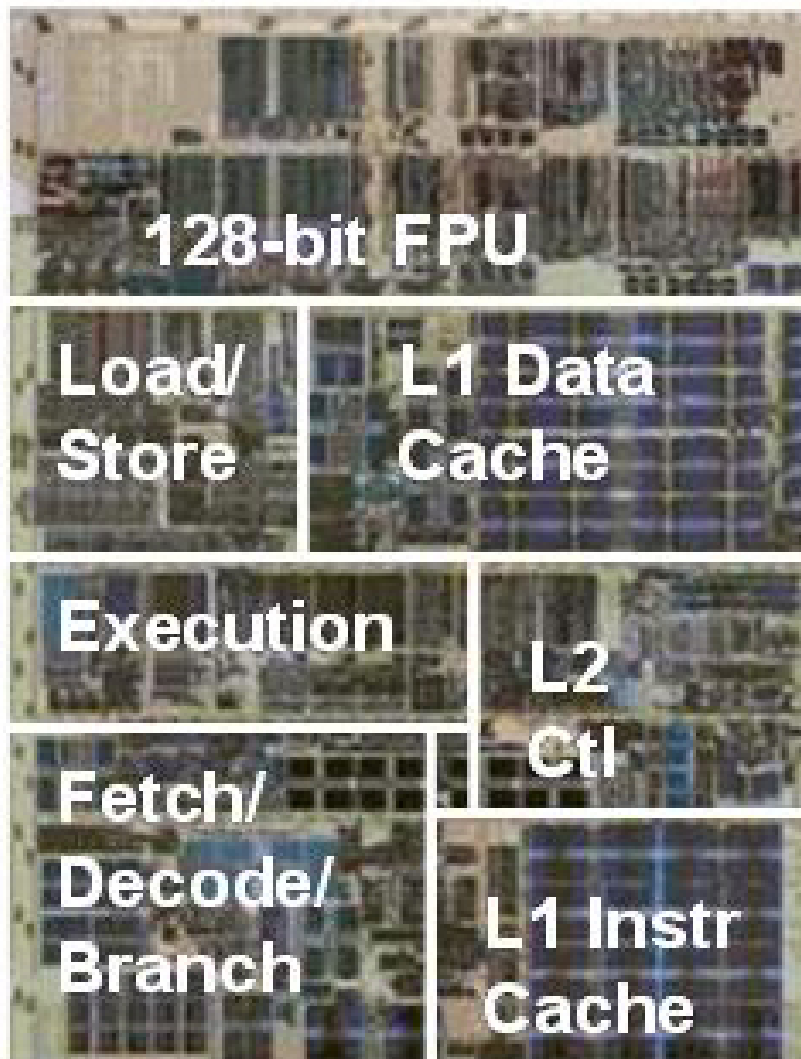
□ PLL

□ Thermal

# Core/L2 Overview

## Core


## L2



# Voltage/Clock Domains

Multiple supplies for power optimization and isolation

- VDDCORE: 0.8V-1.4V
  - Core and L2: 2.0GHz and up
- VDDNB: 0.8V-1.4V
  - Northbridge and L3: 75% of core
- VLDT: 1.2V
  - HyperTransport links
- VDDIO: 1.8V (VTT:0.9V)
  - DDR I/O
- VDDA: 2.5V
  - PLLs (10 across the die) + Thermal



Power and Clock Domains

- VDDCORE Blocks: Core, L2
- VDDNB Blocks: Northbridge, L3
- PLL
- VLDT Blocks: Hypertransport
- VDDIO Blocks: DDR I/O
- ▲ ck Differential Clock Receiver

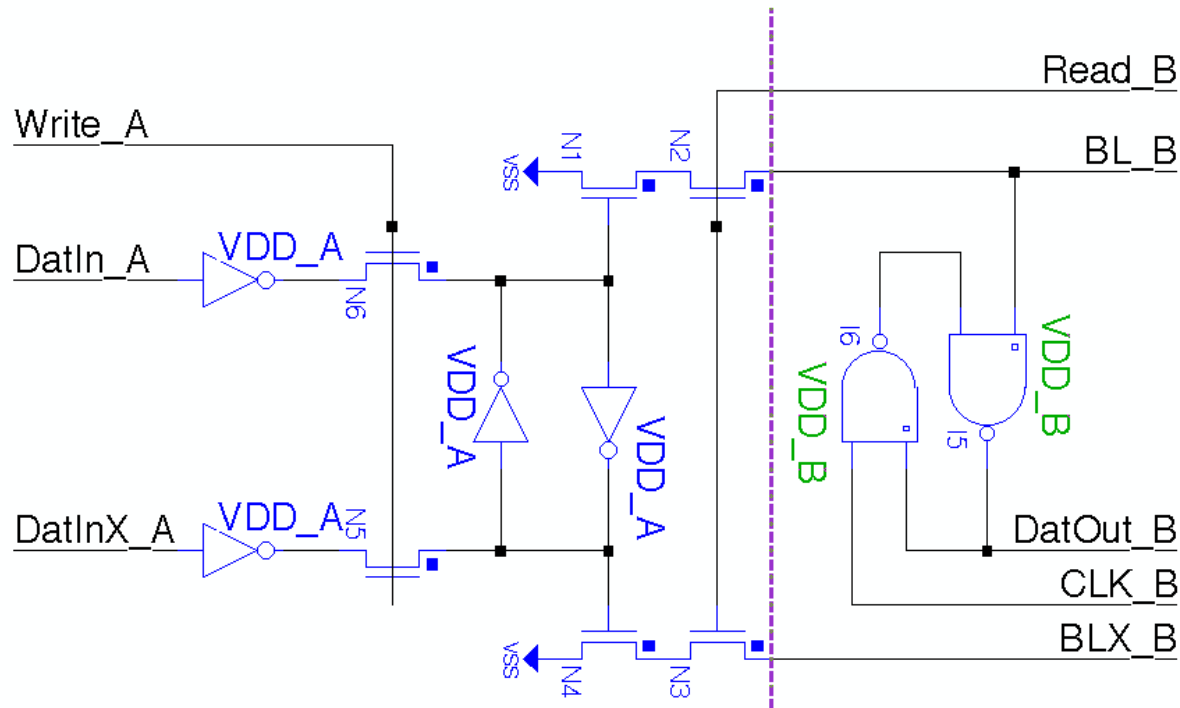


# Clock Domains

- Regional clock domains: lower power, tighter skew (12ps)
- Per core, independent PLL running at 2+GHz
- Independent Northbridge (NB) PLL running at  $\frac{3}{4}$ \*core
- L3 uses the NB clock but at arbitrary phase
- DDR and HT interfaces have independent PLLs
- Source Synchronous clocked Buses connect I/O to NB.
- NB-2-core & core-2-NB communication: 16 entry FIFO.
  1. Asynchronous : Max performance: ~2 entry pointer separation.
  2. Synchronous: Provides deterministic operation for ATE.

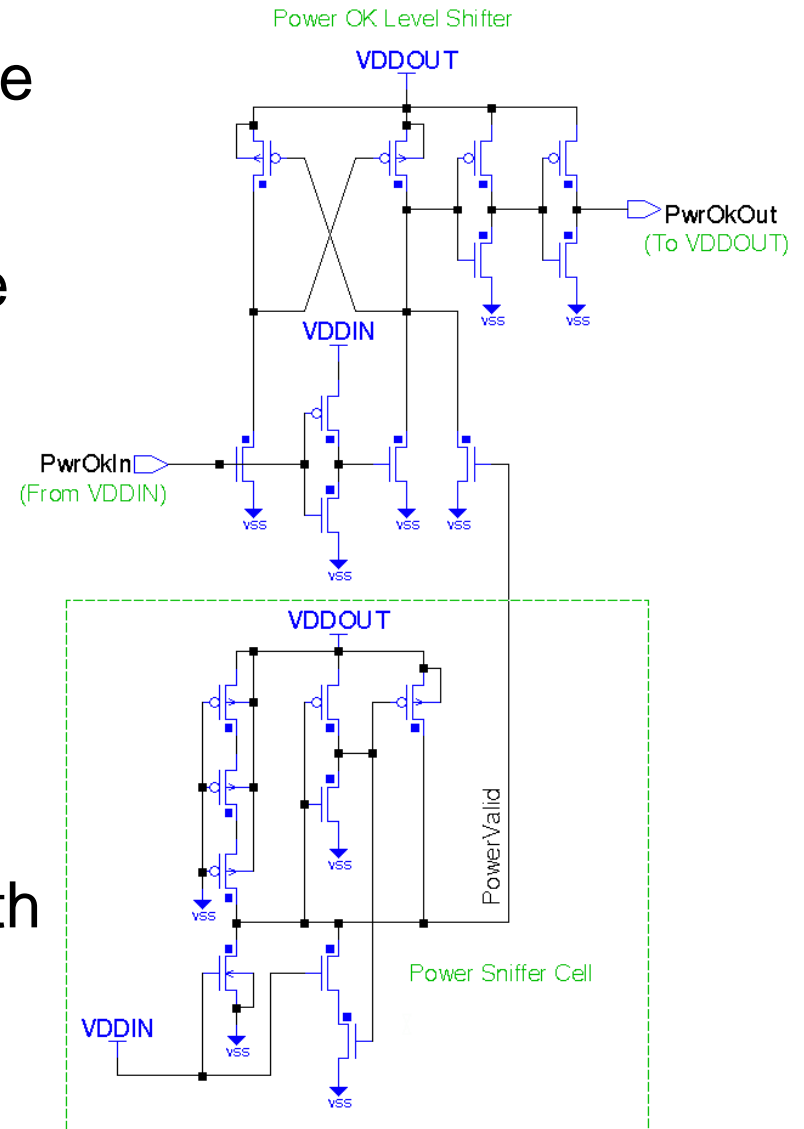
# Domain Crossing

- Domains have: different PLLs (wander) and voltages
- Wide operating range: 0.7 to 1.5V on either side
- Improved latency by embedding into a dynamic FIFO
- Precharge to VDD\_B, evaluate with VDD\_A



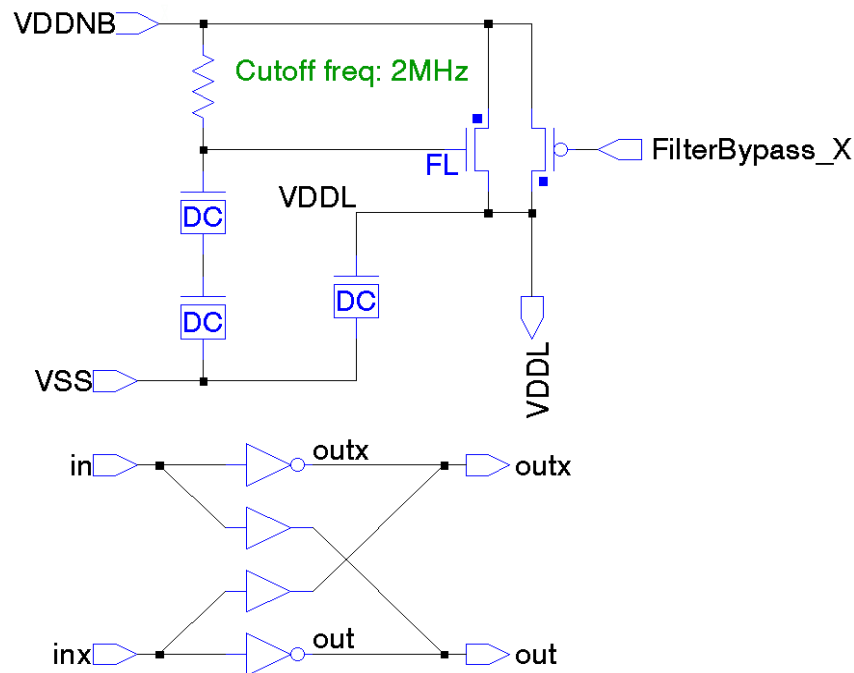
# Power OK Distribution

- PowerOK asserting: Valid voltage to all domains is present
- PowerOK signal originates in the VDDIO domain of the I/Os
- Propagates around the chip in VDDNB
- PowerOK passes from NB to all domains
- Power sniffer cell is designed with hysteresis to prevent spurious switching on power up.



# Noise Filtered Ref Clock

- All PLLs run off of one 200MHz Reference Clock
  - Distributed by a binary tree of specially filtered repeaters
  - Low-pass power-supply filter, 2Mhz pole
  - Reduces feed-through jitter at PLL output: higher Fmax
  - 500ps skew target at 0.8V: deterministic test



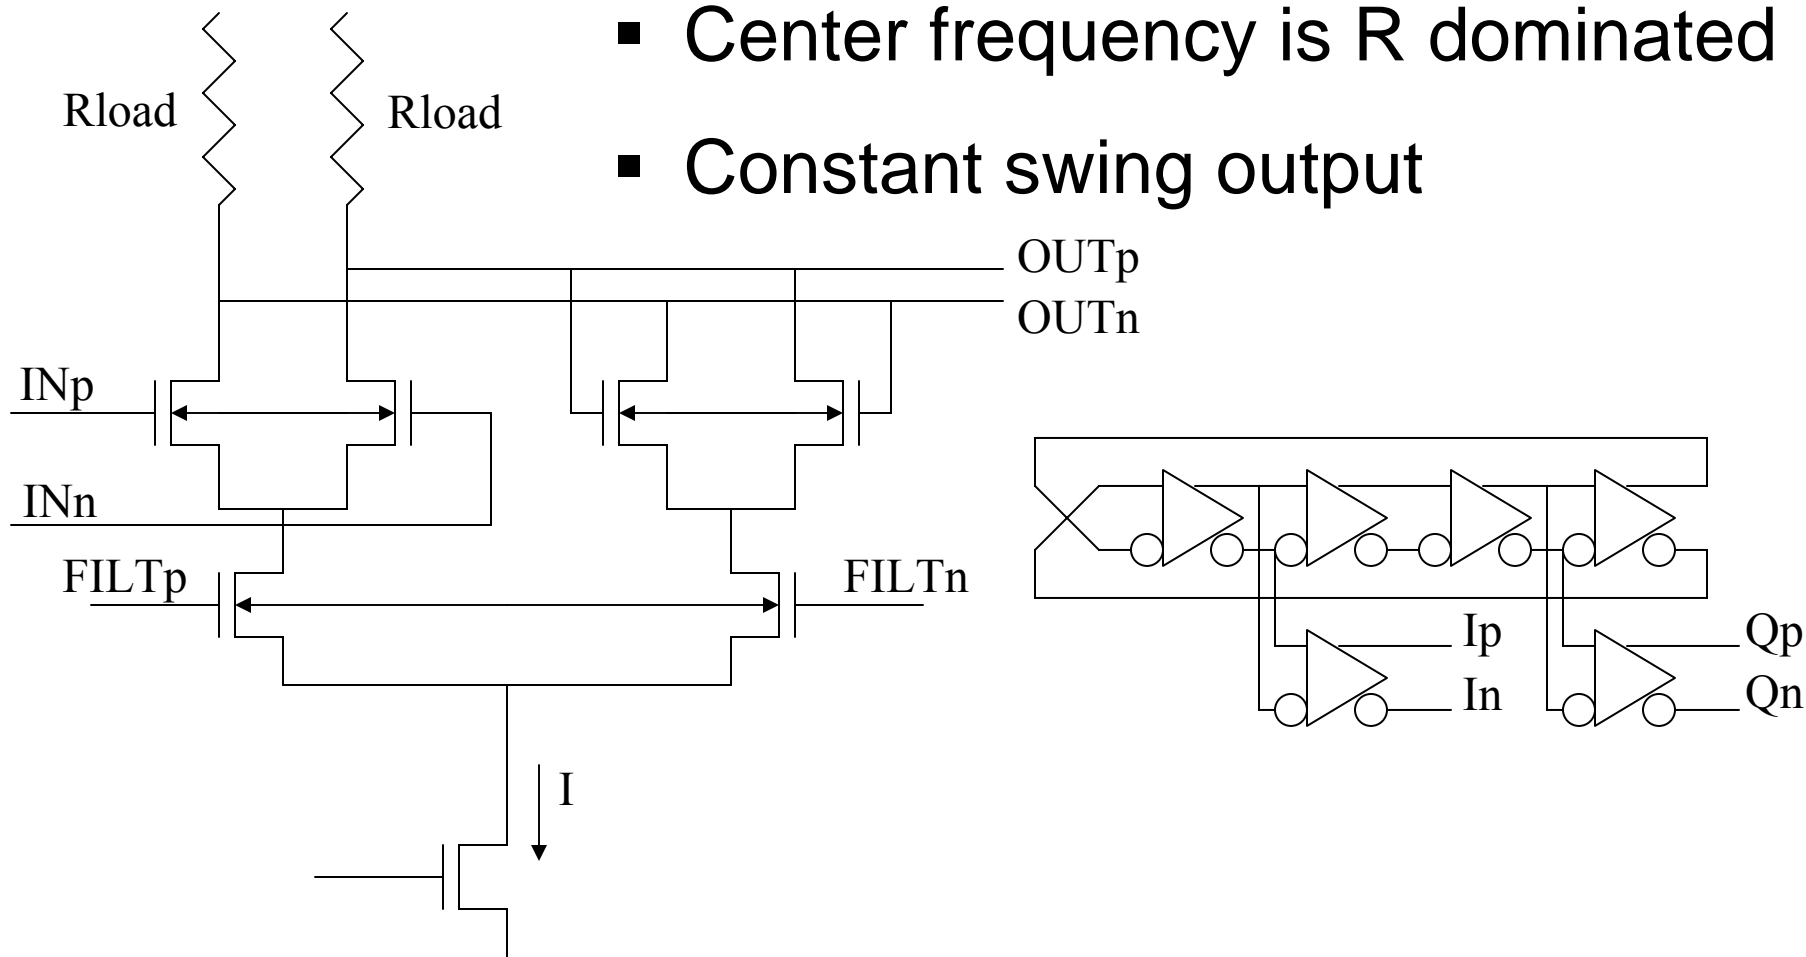
Differential Repeater runs on VDDL supply

# PLL Features

- Accepts 100 or 200 MHz Ref Clock in
- Feedback divider from 1 to 64 (all codes)
- Spine divider (1,2,3,4,6,128,512) post VCO
- Deterministic Spine divider changes (P-state)
- Duty cycle correction on output to spine
- Clock shrink/stretch for speedpath debug

# PLL VCO


- Positive feedback delay control
- Center frequency is R dominated
- Constant swing output



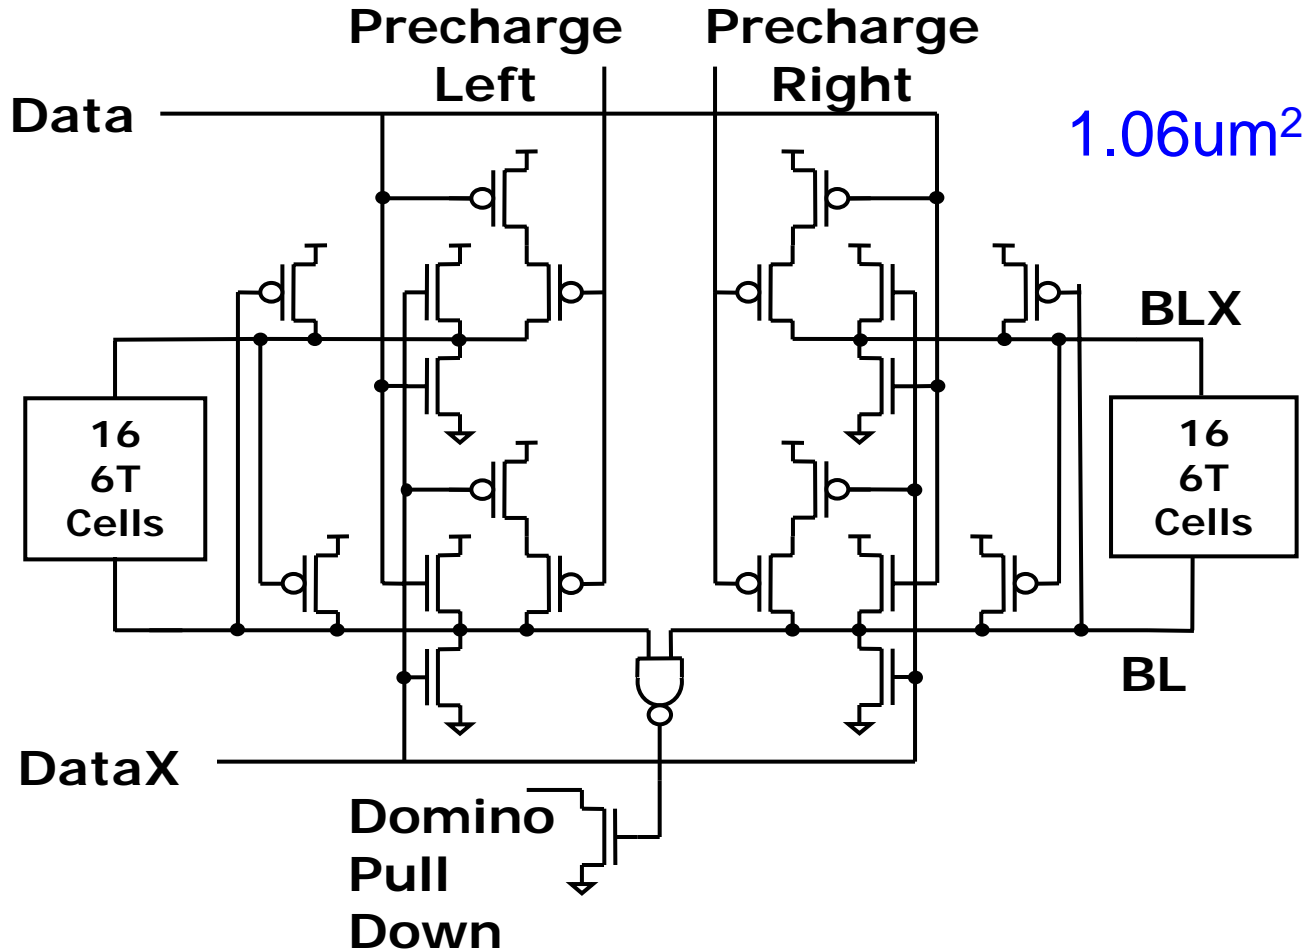


# PLL Closed Loop Response

PLL Close Loop Response Gain, Phase vs Frequency

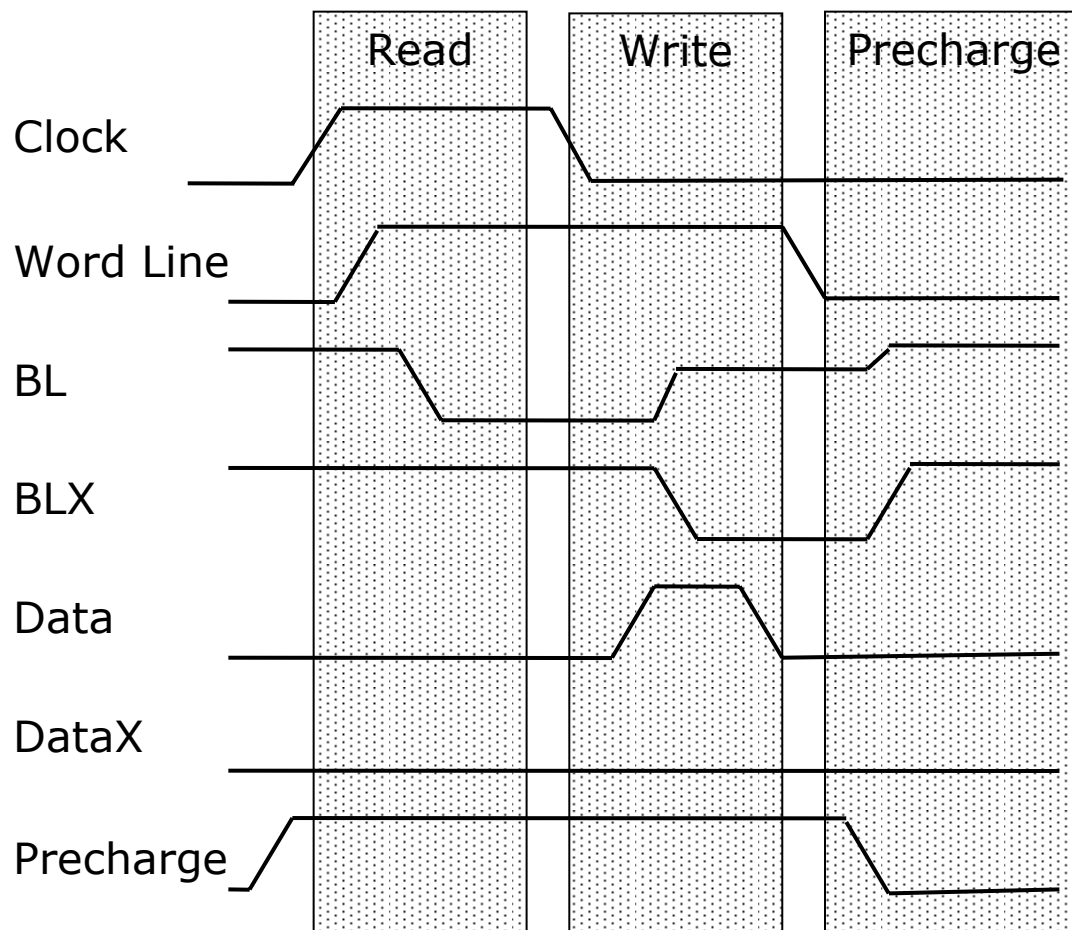


# L1 Cache Bit Slice



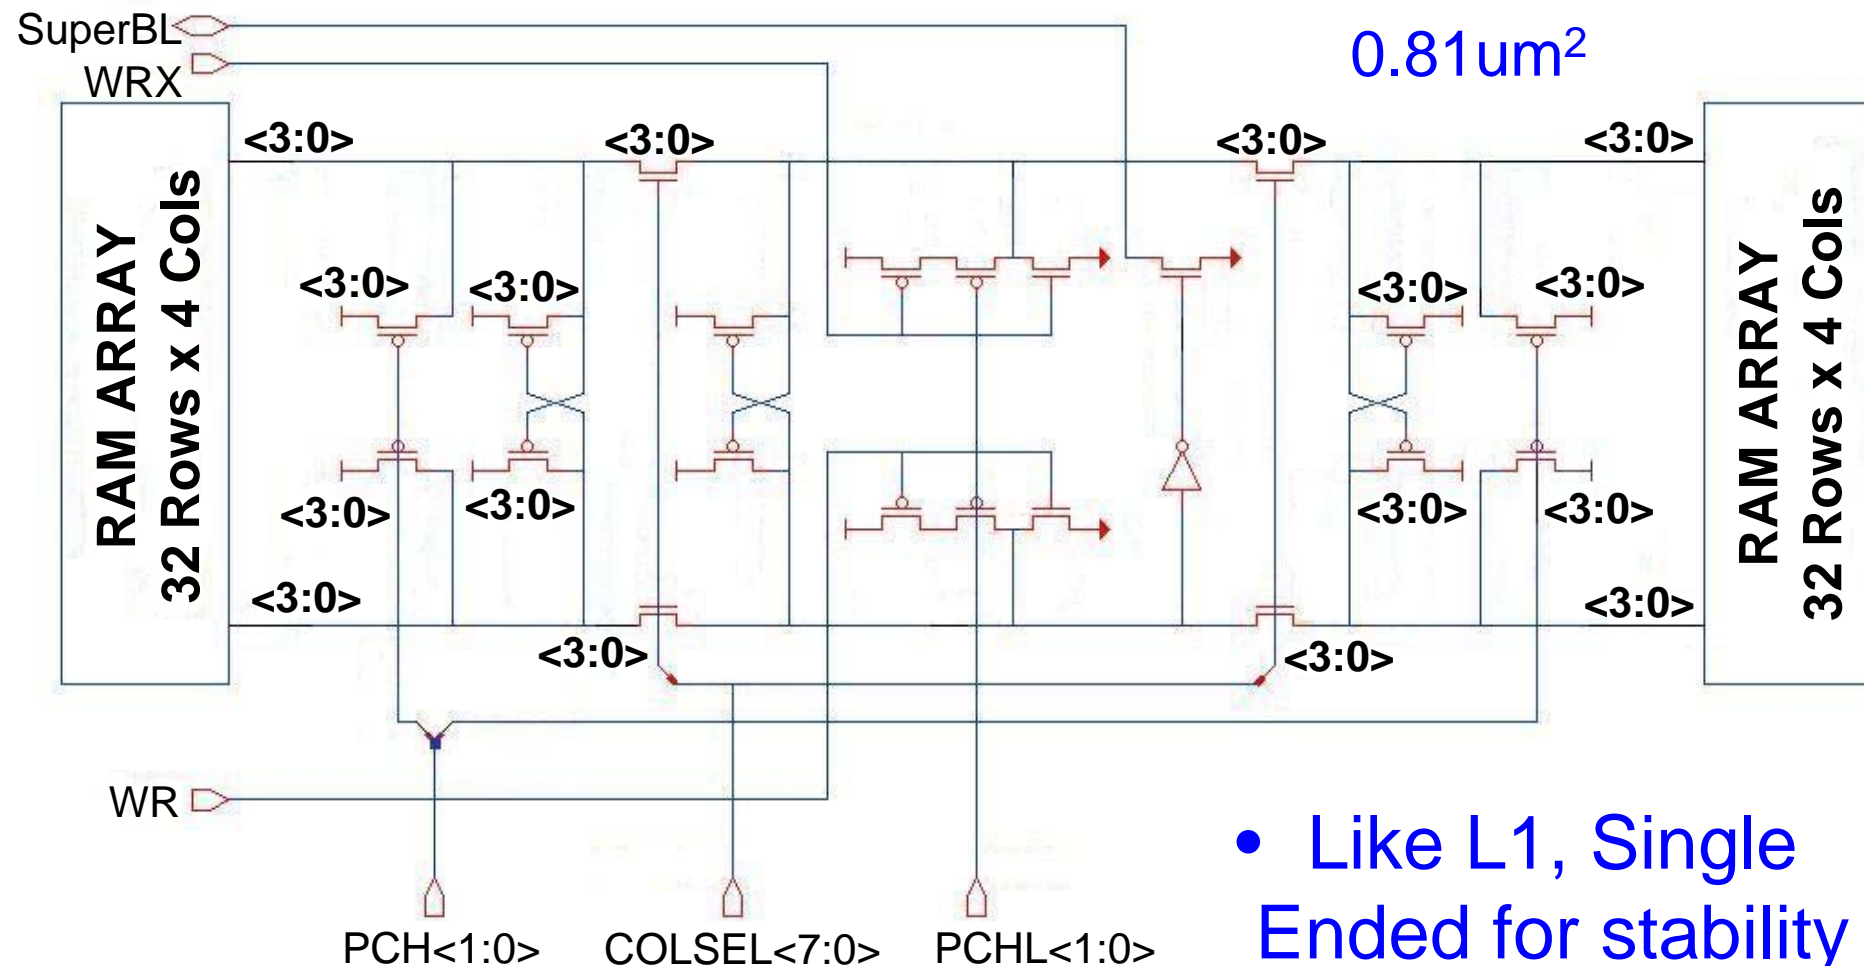
- Single Ended
- Sink and Source for Writes
- Local and Global Bit lines

# L1 Timing: Read Modified Write



- Single Ended for stability
- Precharge is self timed off Write
- Fuse control
- Allows ECC on partial writes


# L2 & L3 Cache Bit Slice



- Like L1, Single Ended for stability
- 3 Level Bit Line

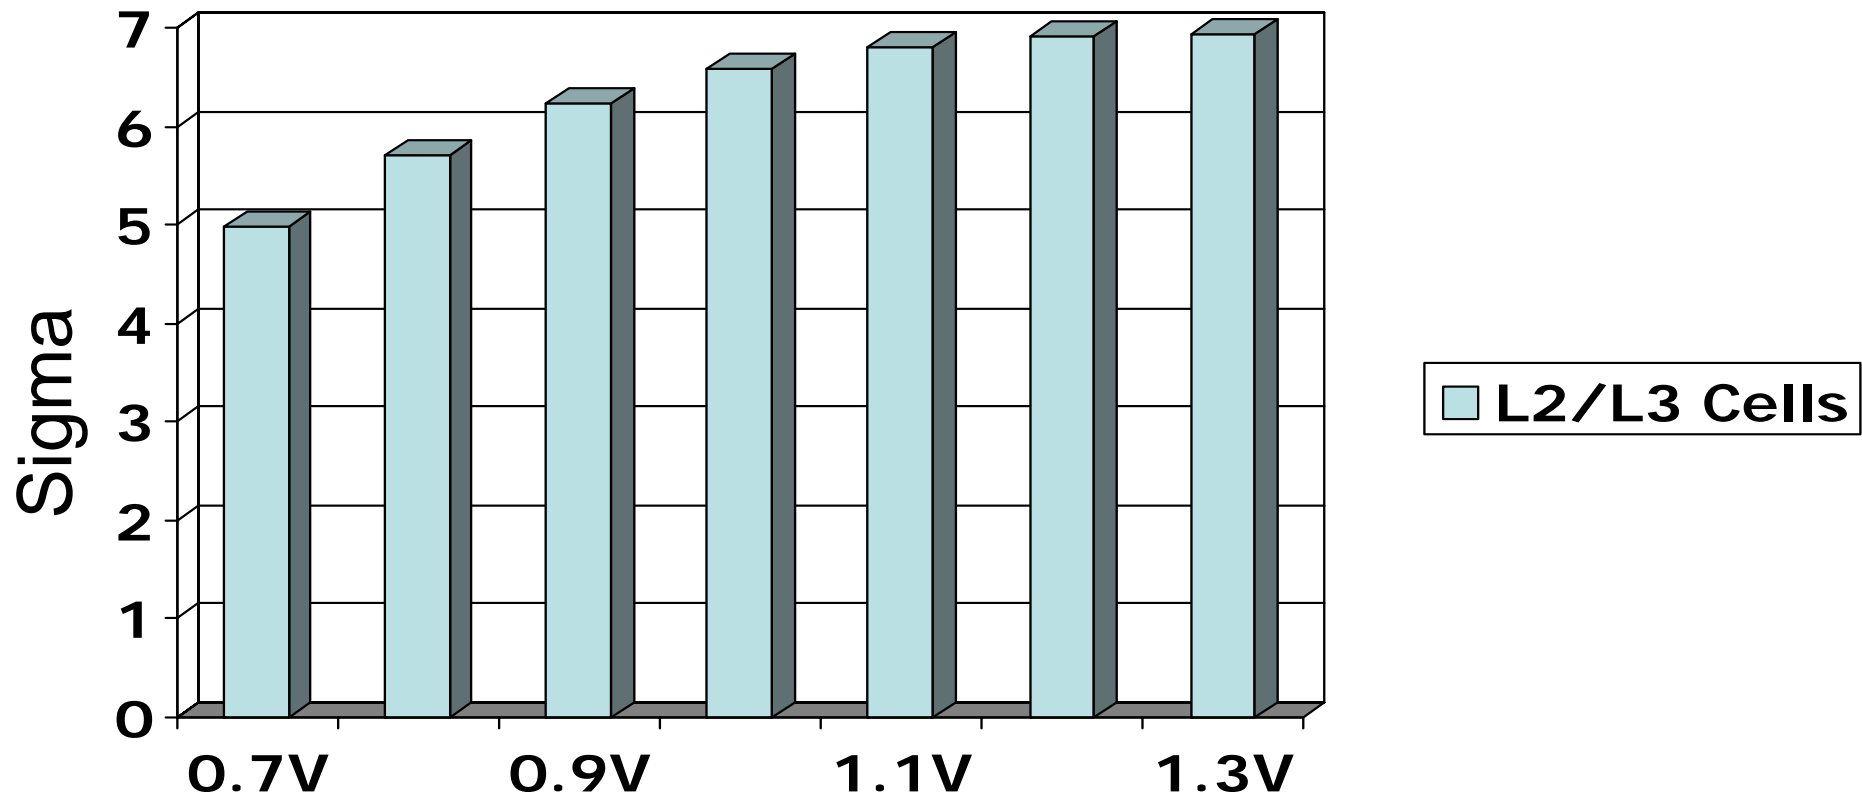
# L3 Timing

- Flexible tiling methodology: liquid L3
- Asynchronous timing for flexible layout



# L2 & L3 Read Disturb Stability

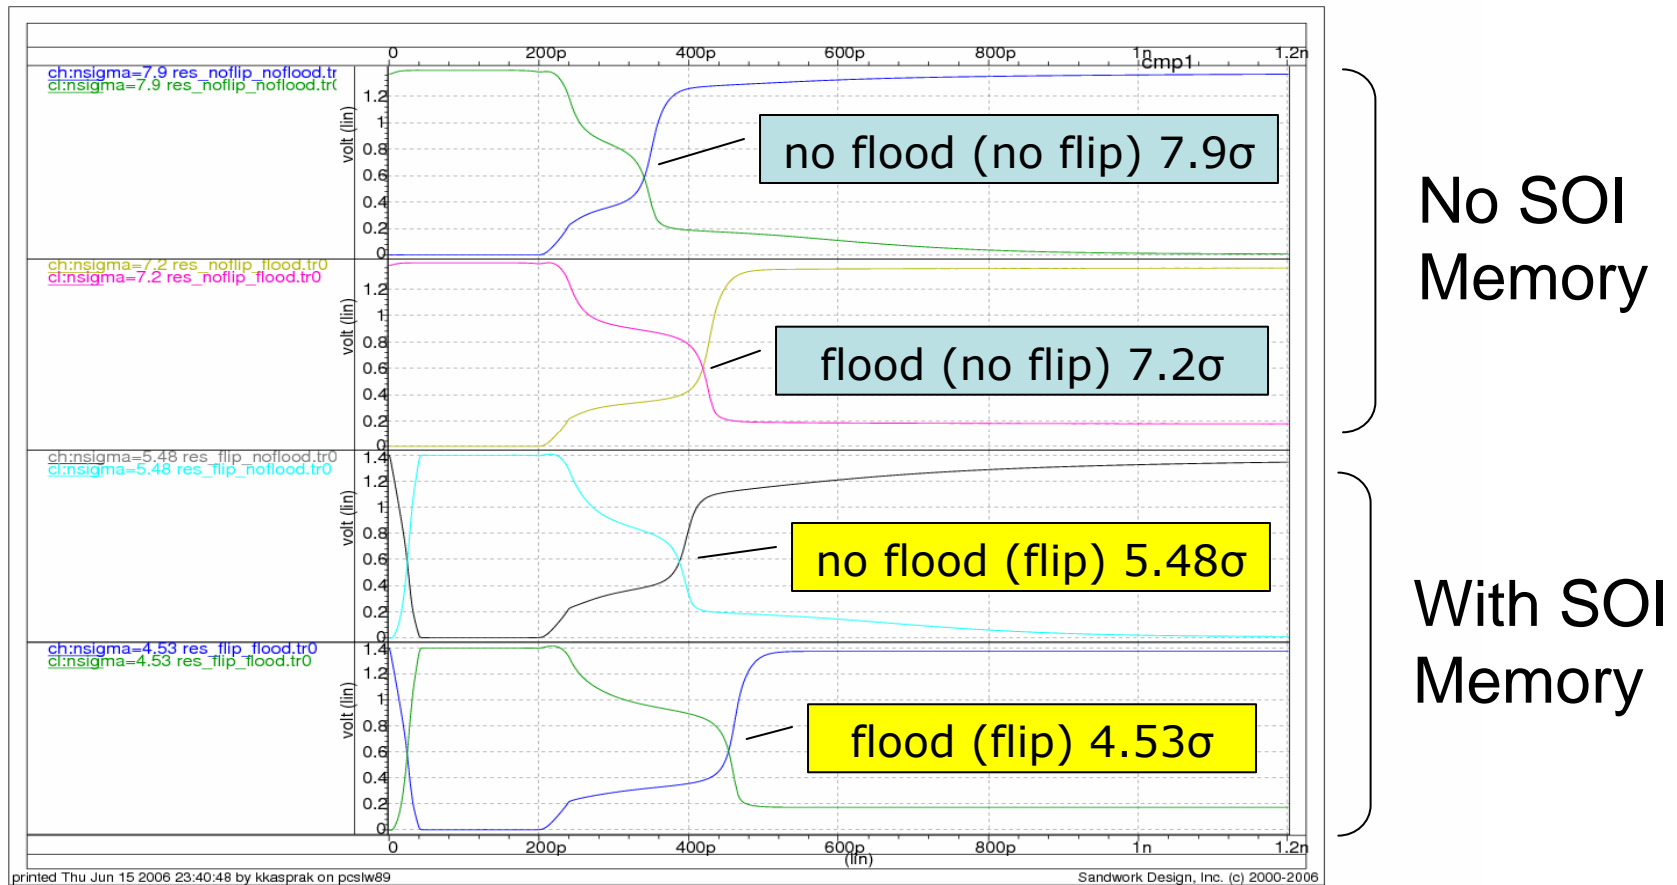
- Product Reliability requires  $> 5$  sigma margin
- Single-ended vs. small-swing: better margin





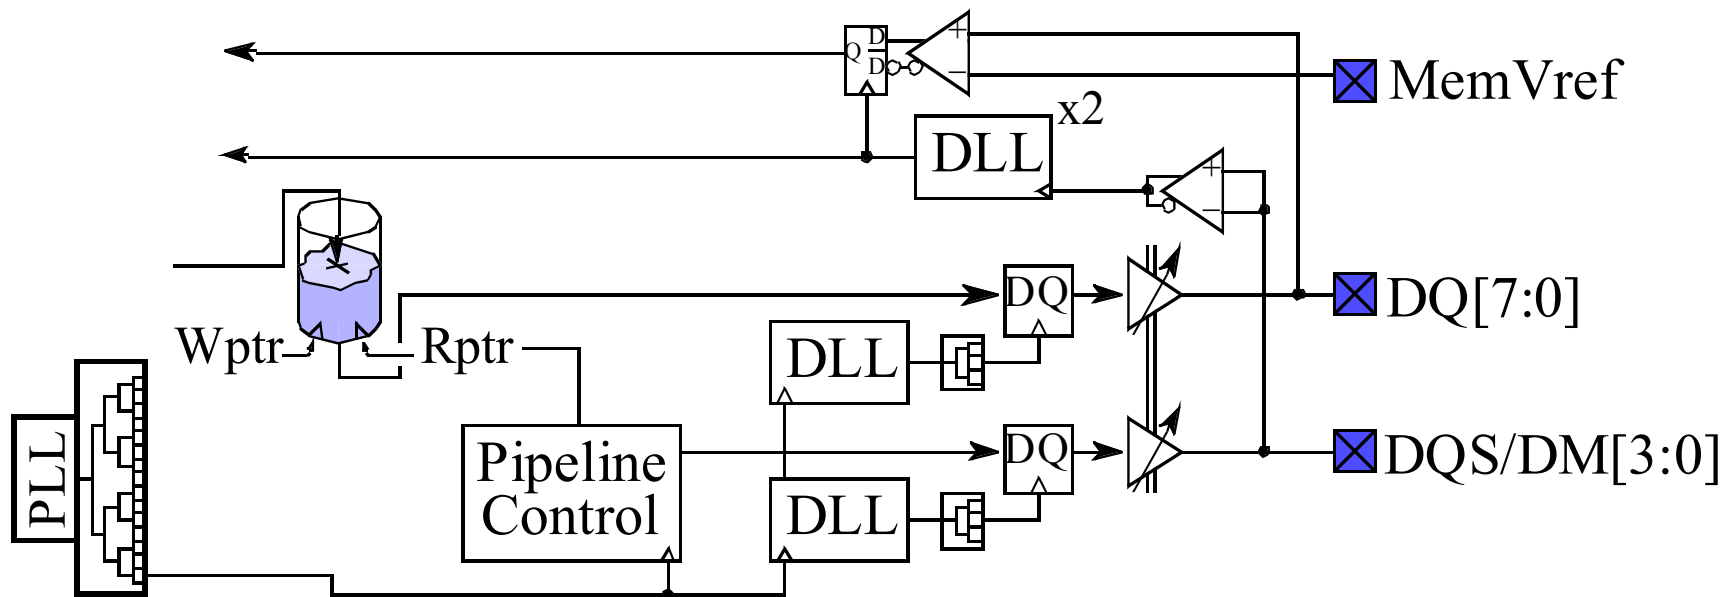
# L2 & L3 Small Swing Study

- Study small swing (sense-amplifier-based)
- Increased  $\Delta V_T$  &  $\Delta L$  variation + SOI Hysteresis

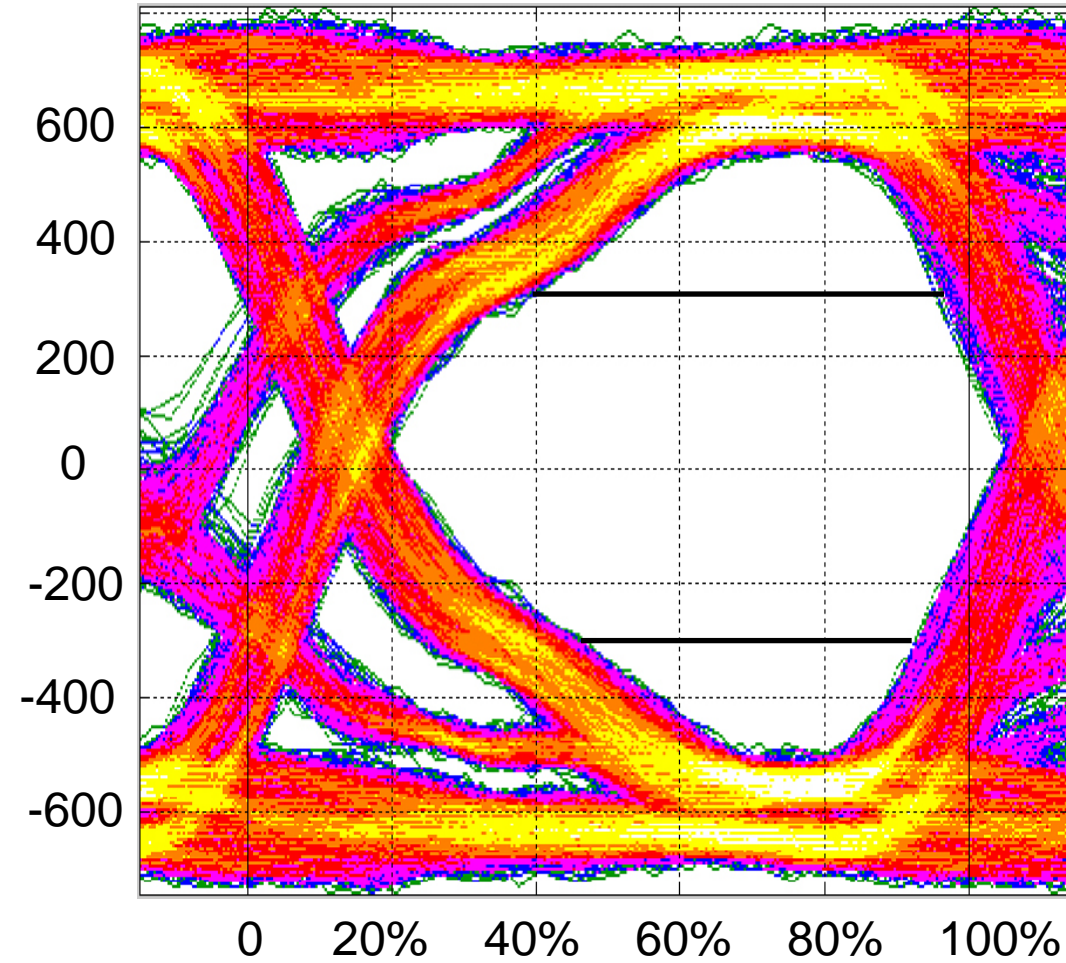


# DDR DRAM I/O

- Supports DDR2 & 3, 1.4 to 1.9V, 400 to 1600Mb/s
- 2 independent (R/W,W/R,R/R,W/W), 64b channels (+ECC)
- 4 DIMM/chan (U,R,SO), per-DIMM trained delay settings.
- 44 DLL/chan with dynamic delay update (per burst).

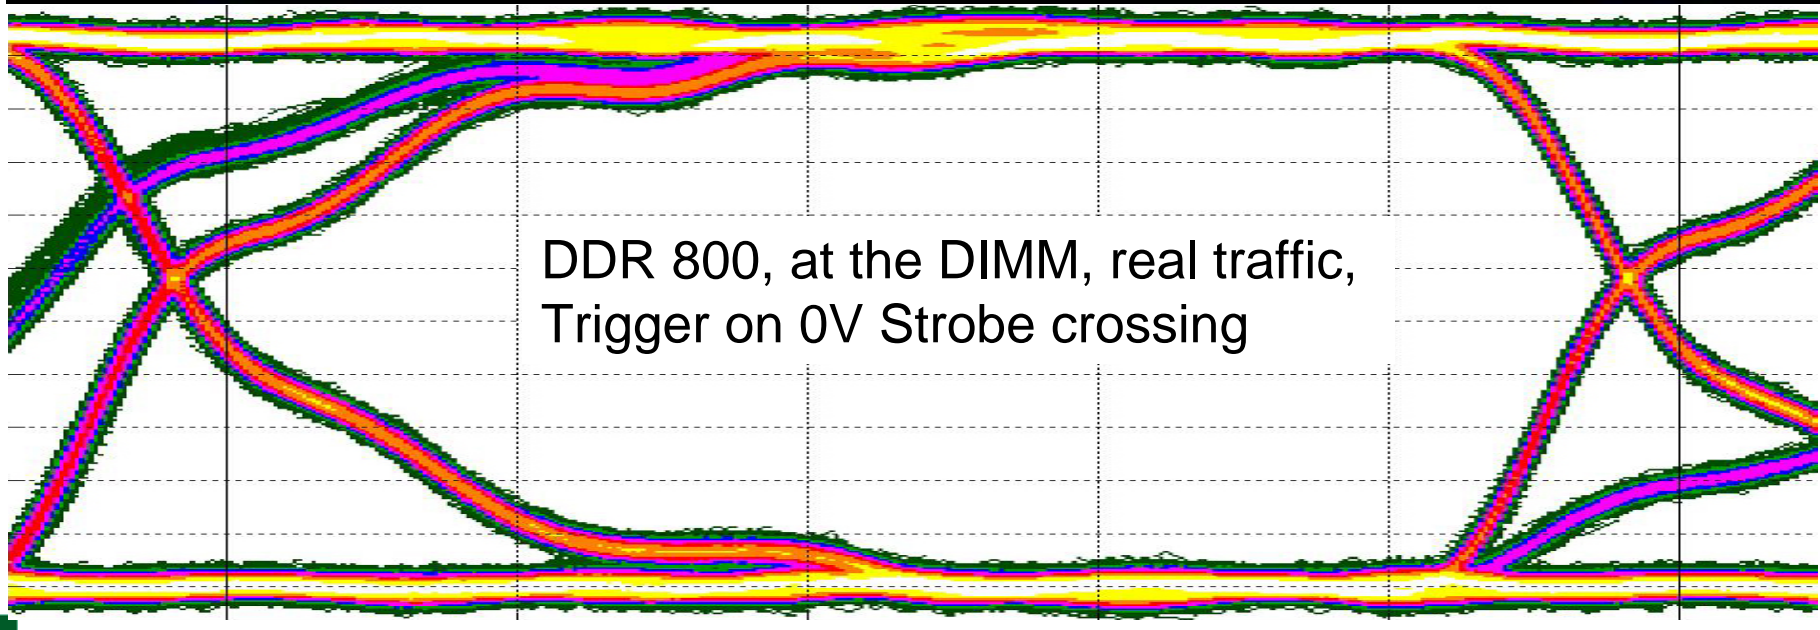



# DDR DRAM Read Eye



- DDR 800 (1250ps UI)
- Real Motherboard
- Under the Socket
- $\pm 300\text{mV}$  margin

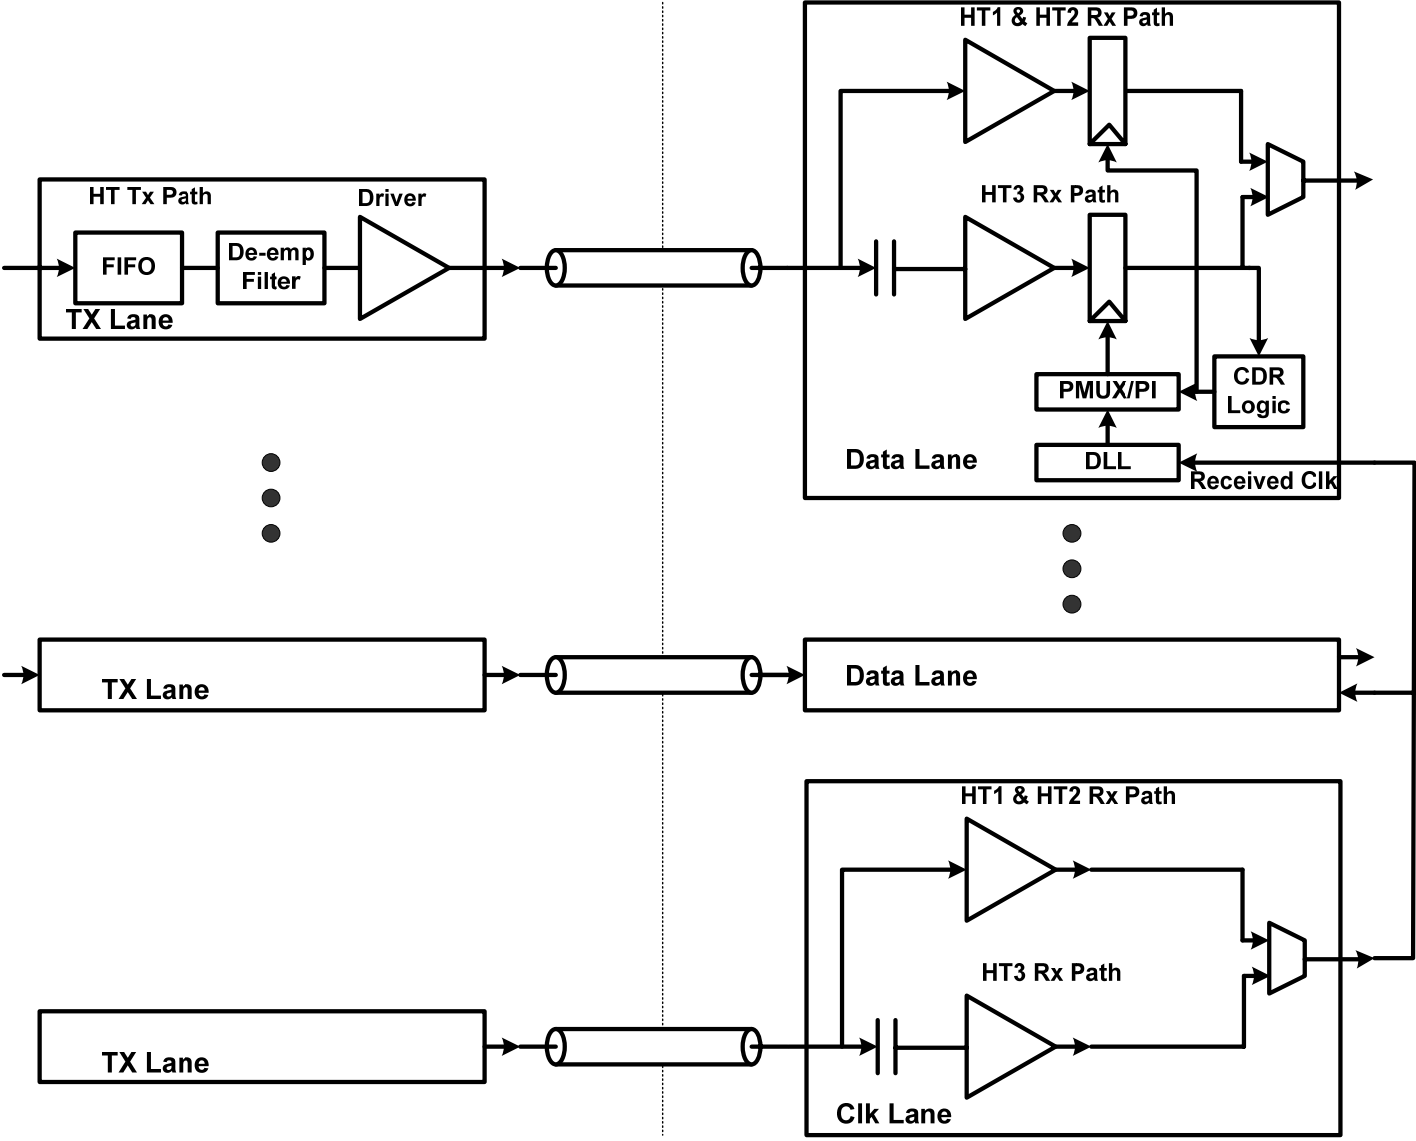
# DDR DRAM Write Eye



# HyperTransport™ I/O

- Backward Compatible (legacy) to HT1
  - 400, 800, 1200, 1600 and 2000 MT/s
- Supports all HT3 speed
  - 2.4, 2.8, 3.2, 4.0, 4.4, 4.8 and 5.2 GT/s
- Support various power saving modes
- Tolerates large HT3 common mode variation
- Linear TX equalization (de-emphasis)
- Non-linear RX equalization (DFE)


# HT Block Diagram



- Clock forwarded
- Simple DLL CDR
- SSC tolerant

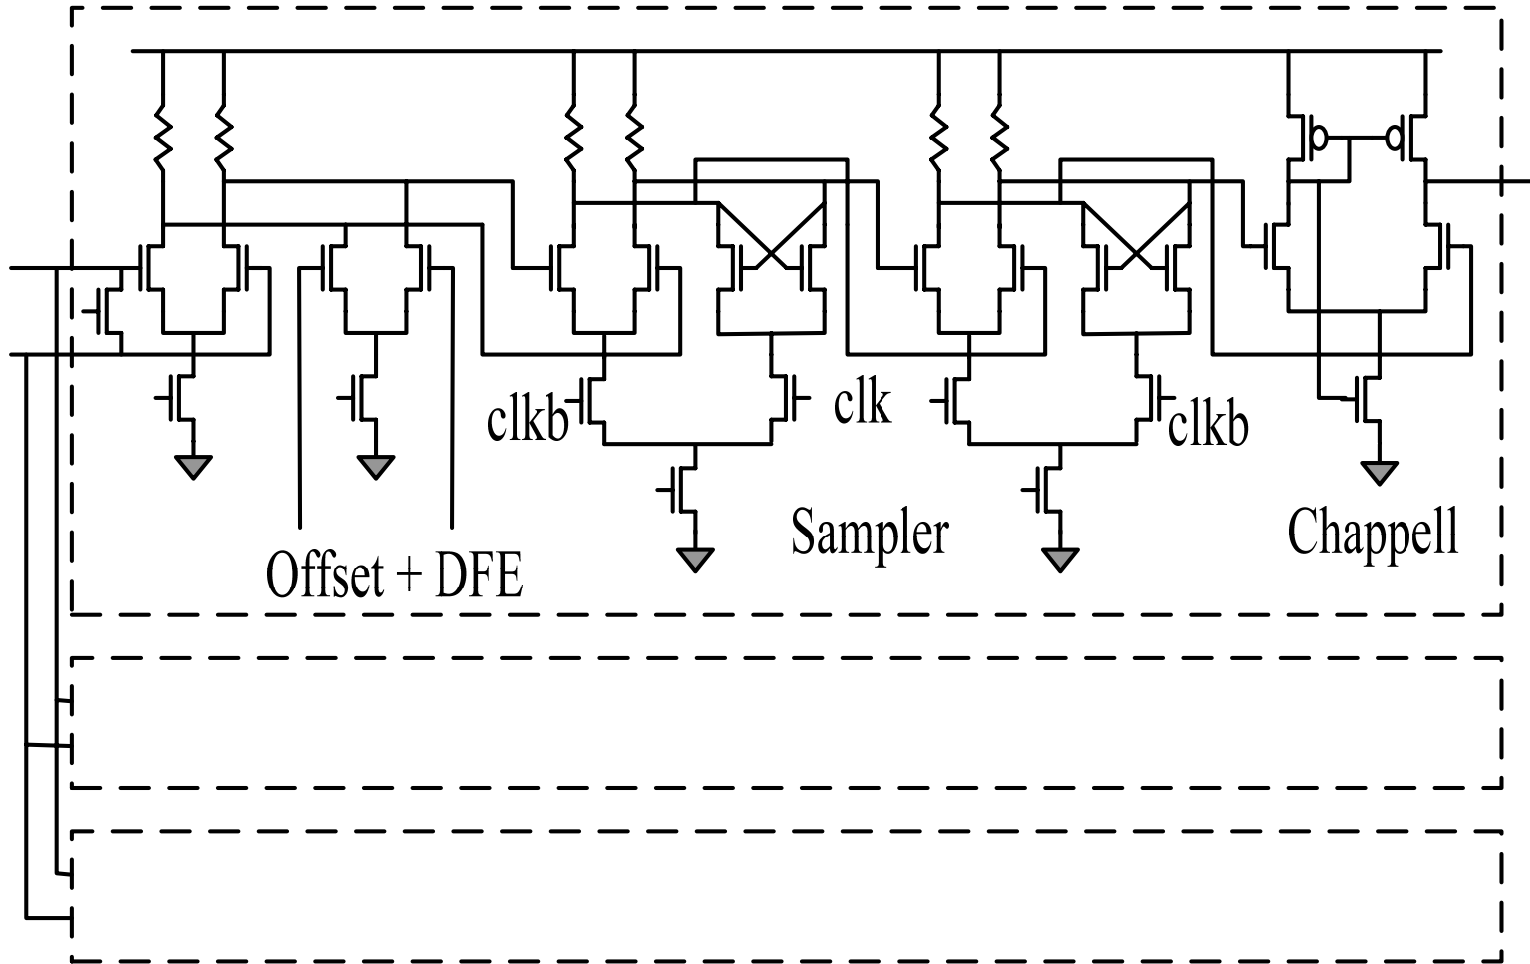


# HT Driver with Linear Eq

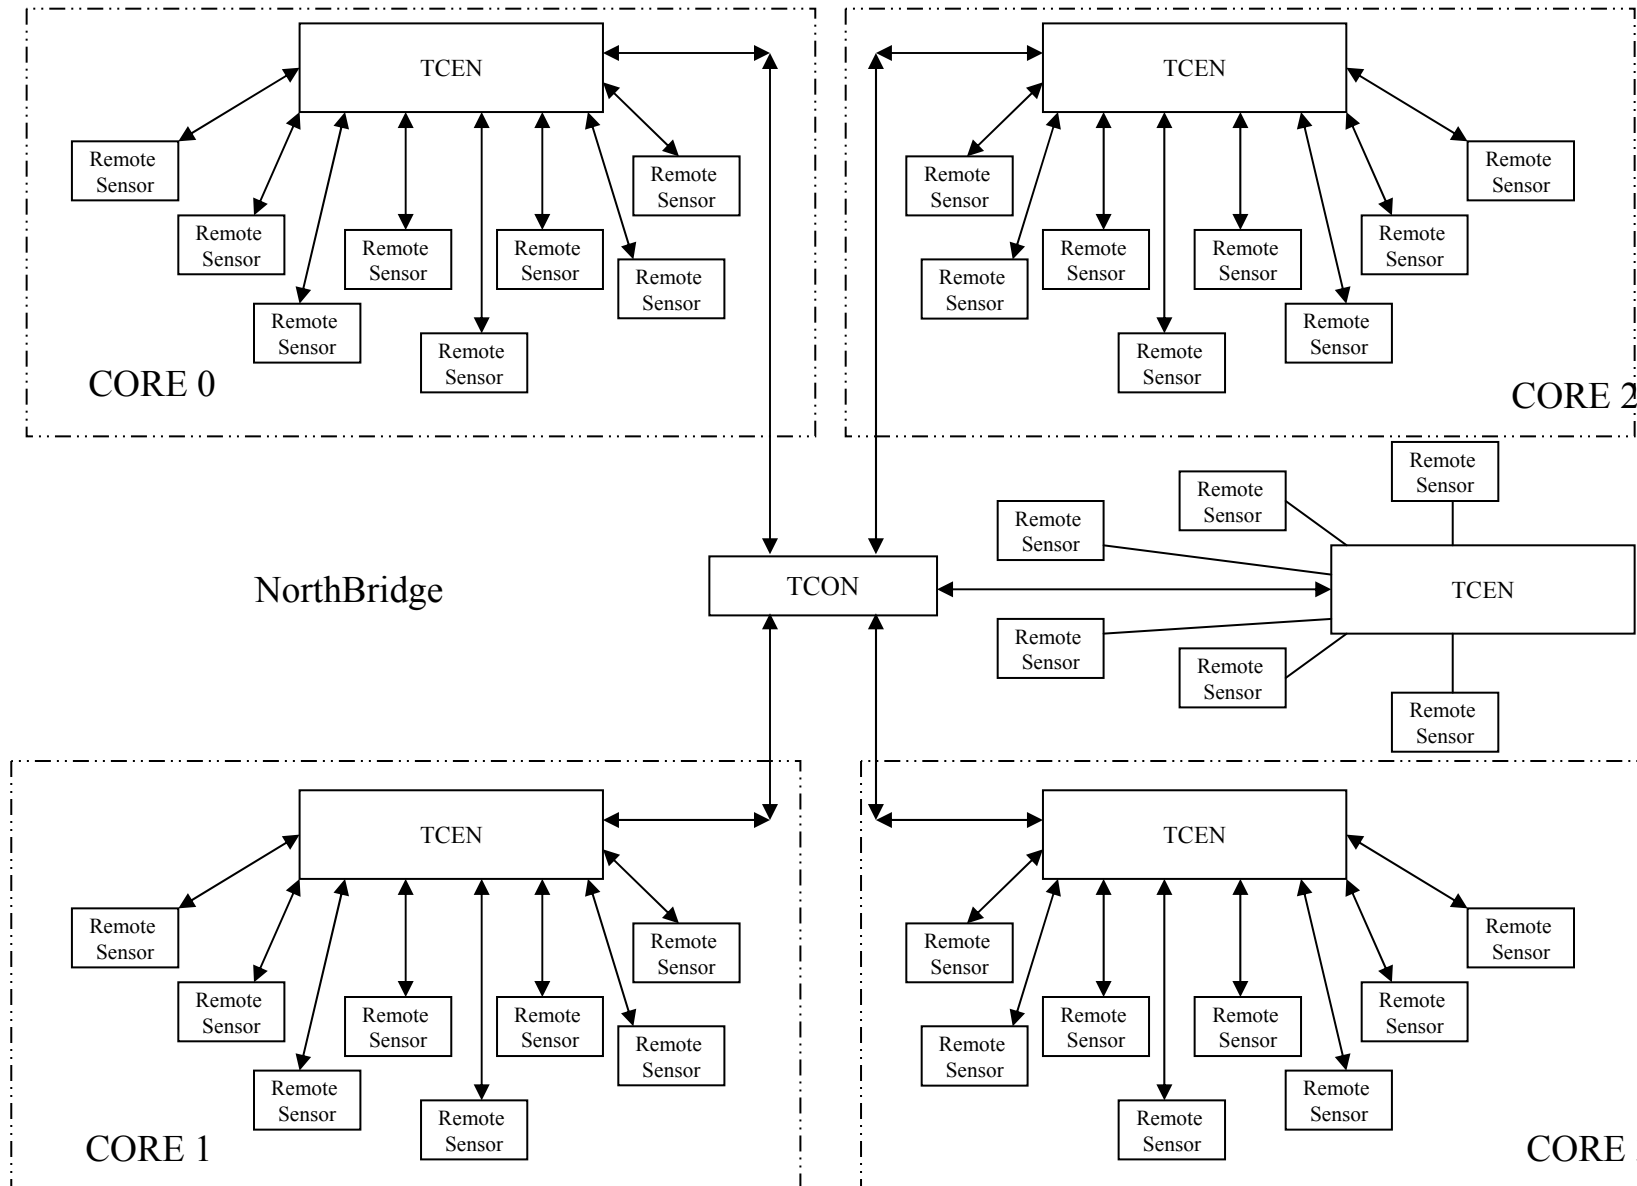


- Primary Voltage mode Driver
- Eq done with current sources
- Lower power when  $> 3\text{dB}$  of boost

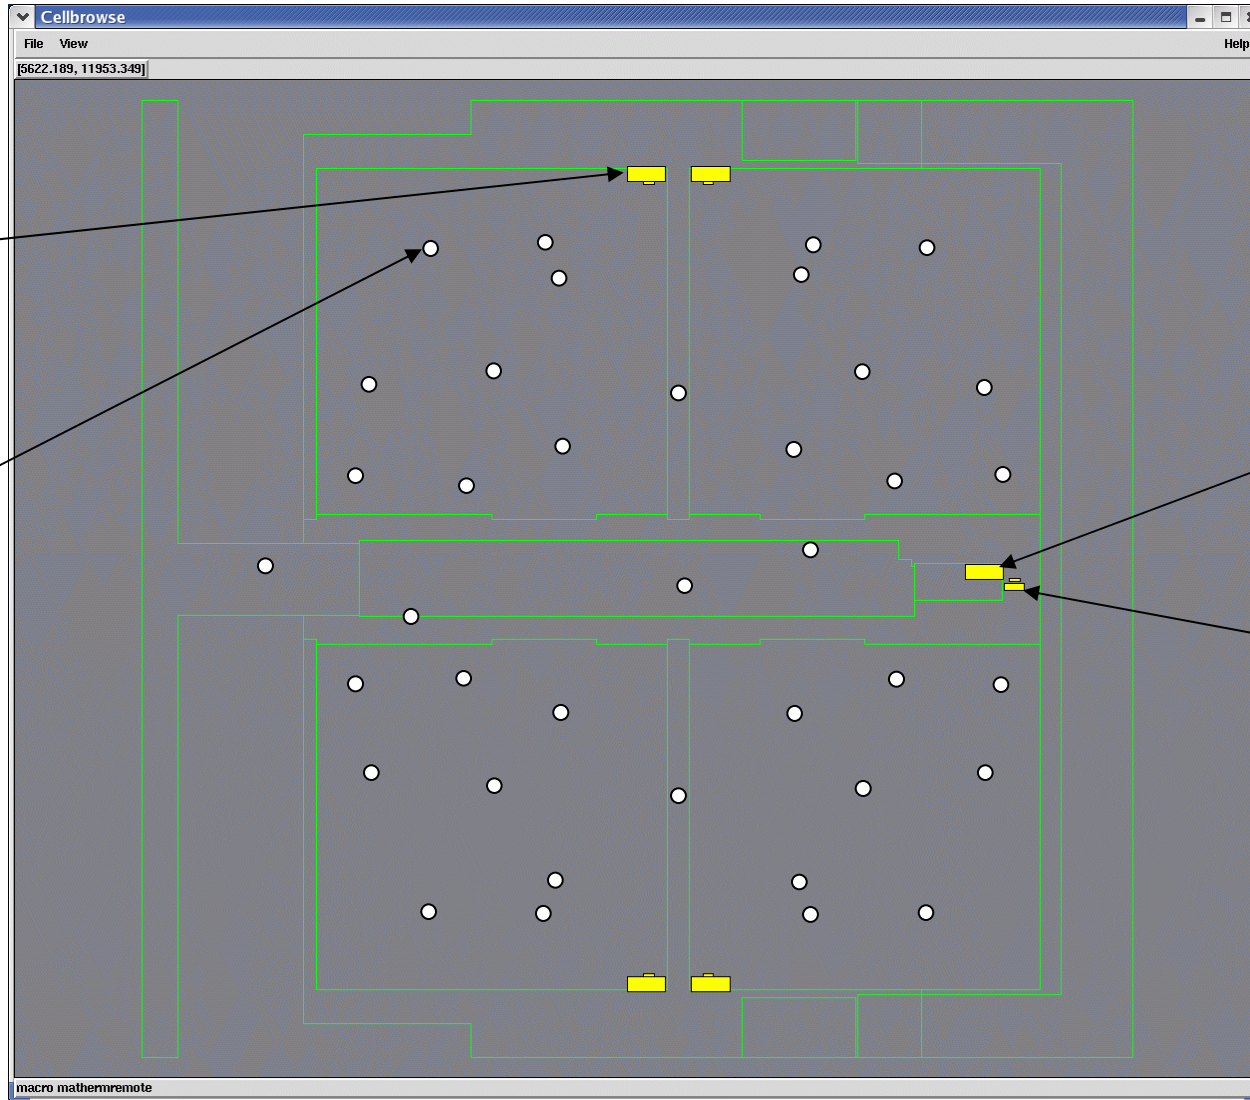
# HT RX with DFE



# Thermal: Chip Wide Block Diagram



# Thermal Sensor Locations



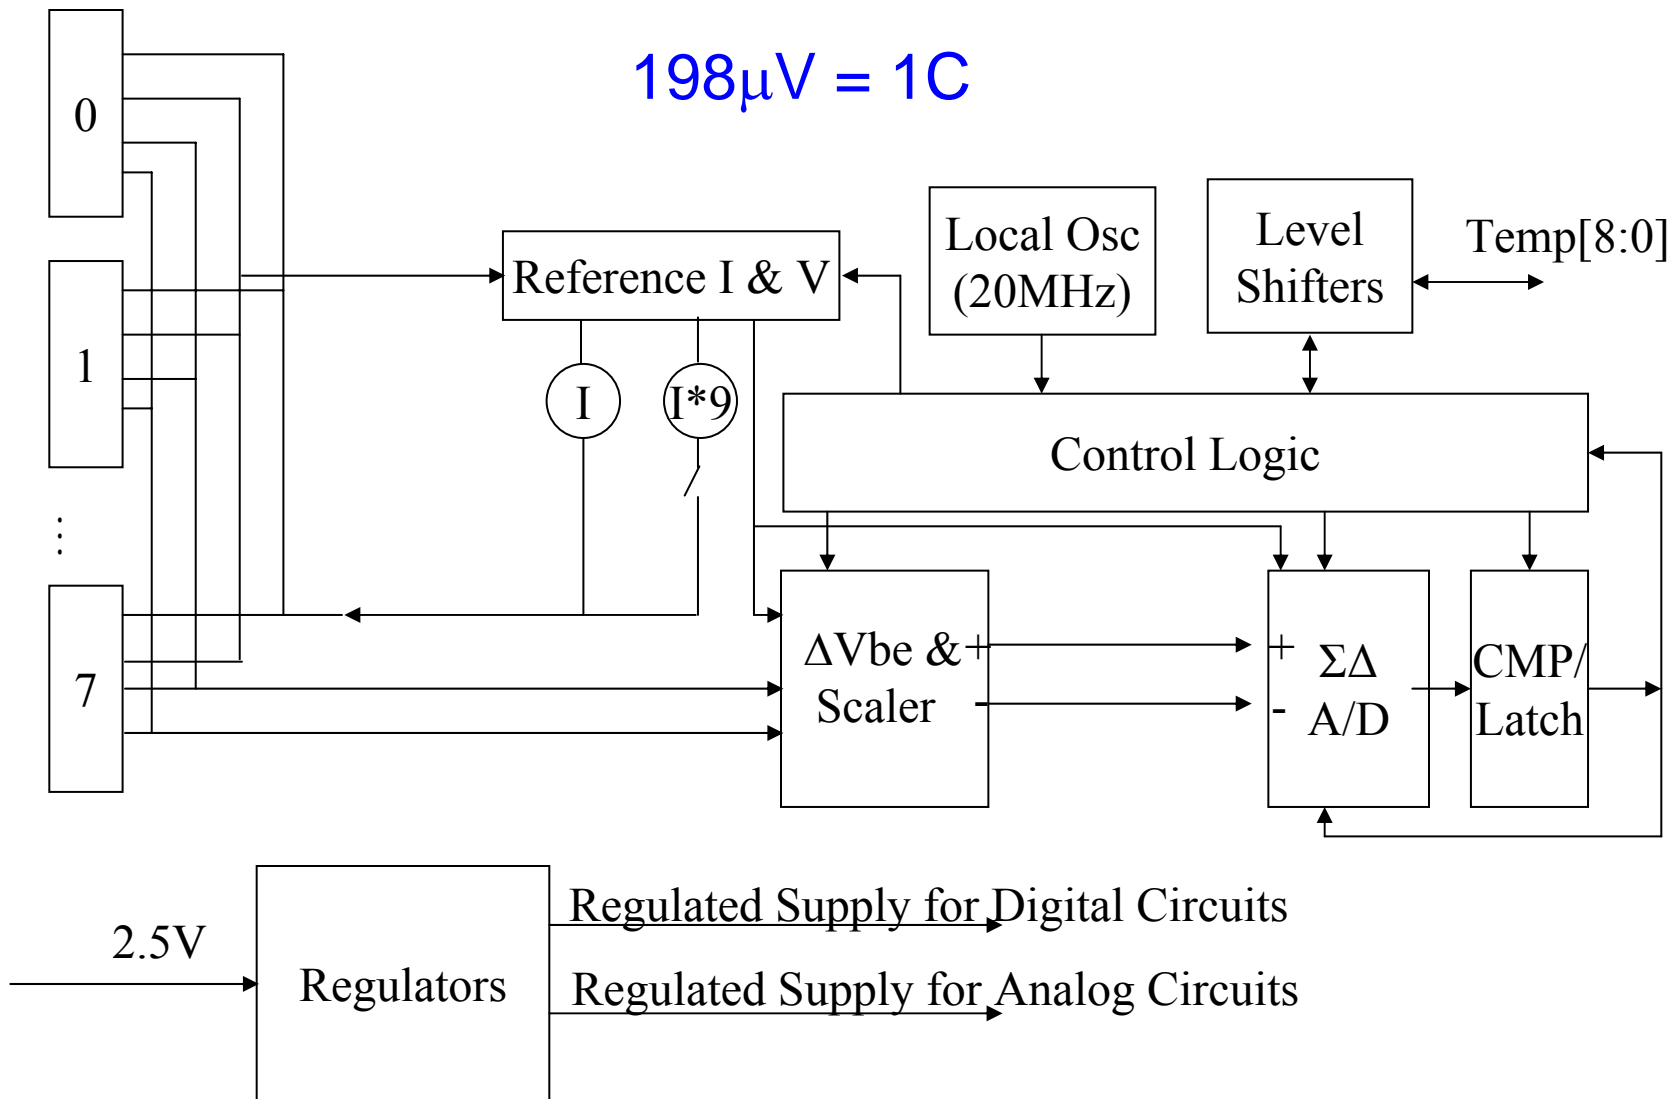
ThermCenter

ThermCenter

Remote  
Diode  
Sensor

External  
Diode

# Thermal: Block Diagram



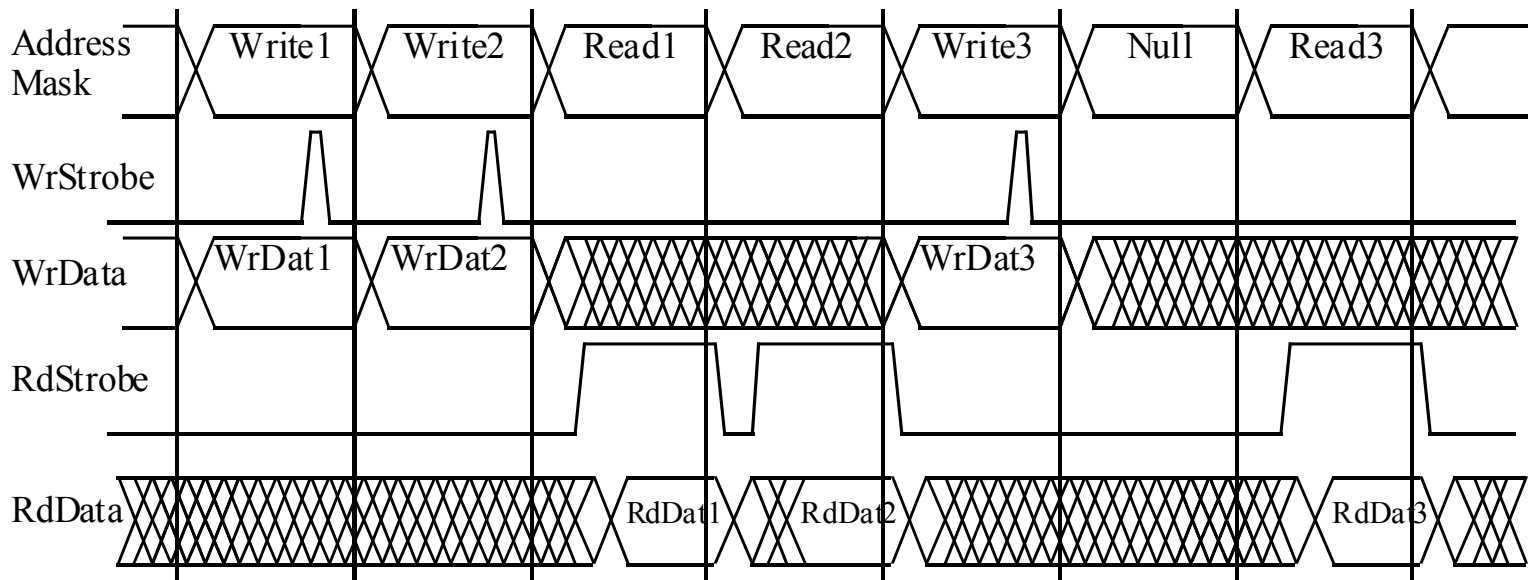
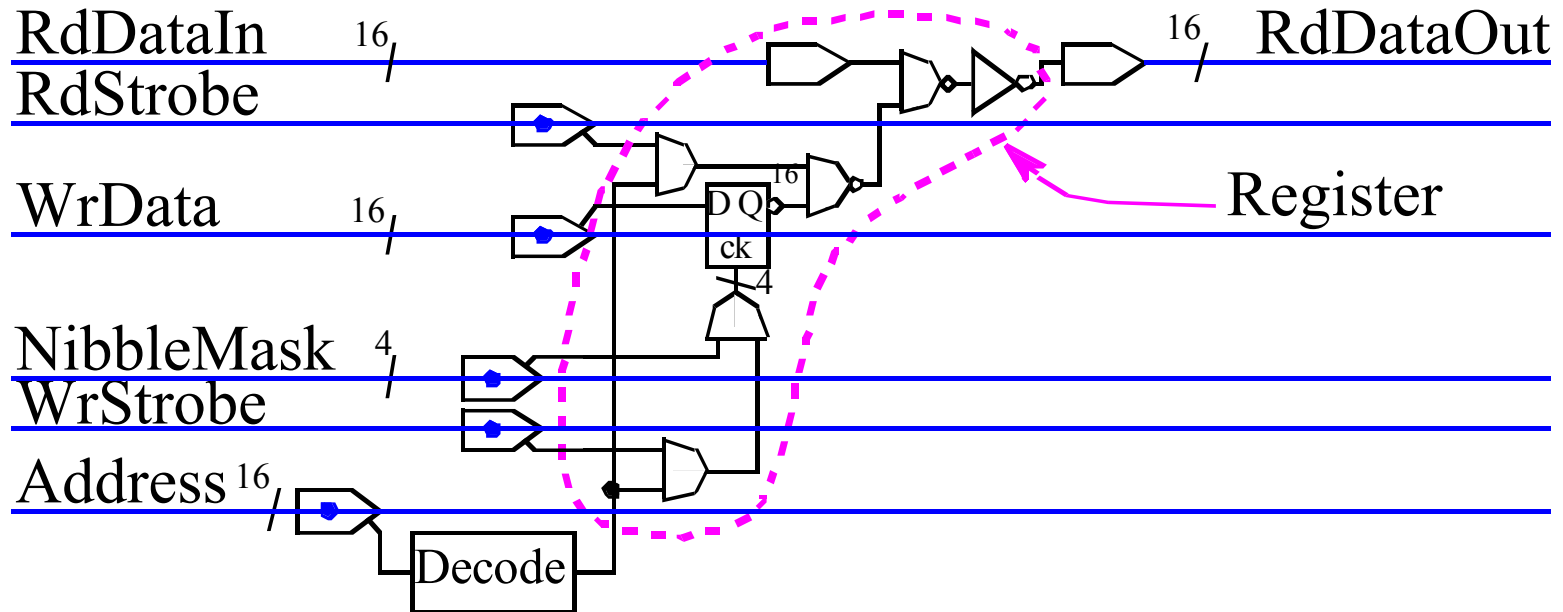
# Design Study (DS): DDR Phy

- Advice for Mixed-Signal IC designers
- Technology does not care about you
- When you overload functions, beware
- Always have software hooks (use all modes)
- Poly (or other) fuses are your friends
- Put Defaults in metal (one layer to change)
- Use servo loops to set bias conditions
- Use DACs for the servo loop reference

# DS DDR: Register Space

- Use scalar bits if you have less than 50 bits
- Otherwise, use a simple bus and register
- Physical locality to where the bits are used
- Extensible: Easy to add bits and functions
- Add in write masks and broadcasts
- It will save your bacon

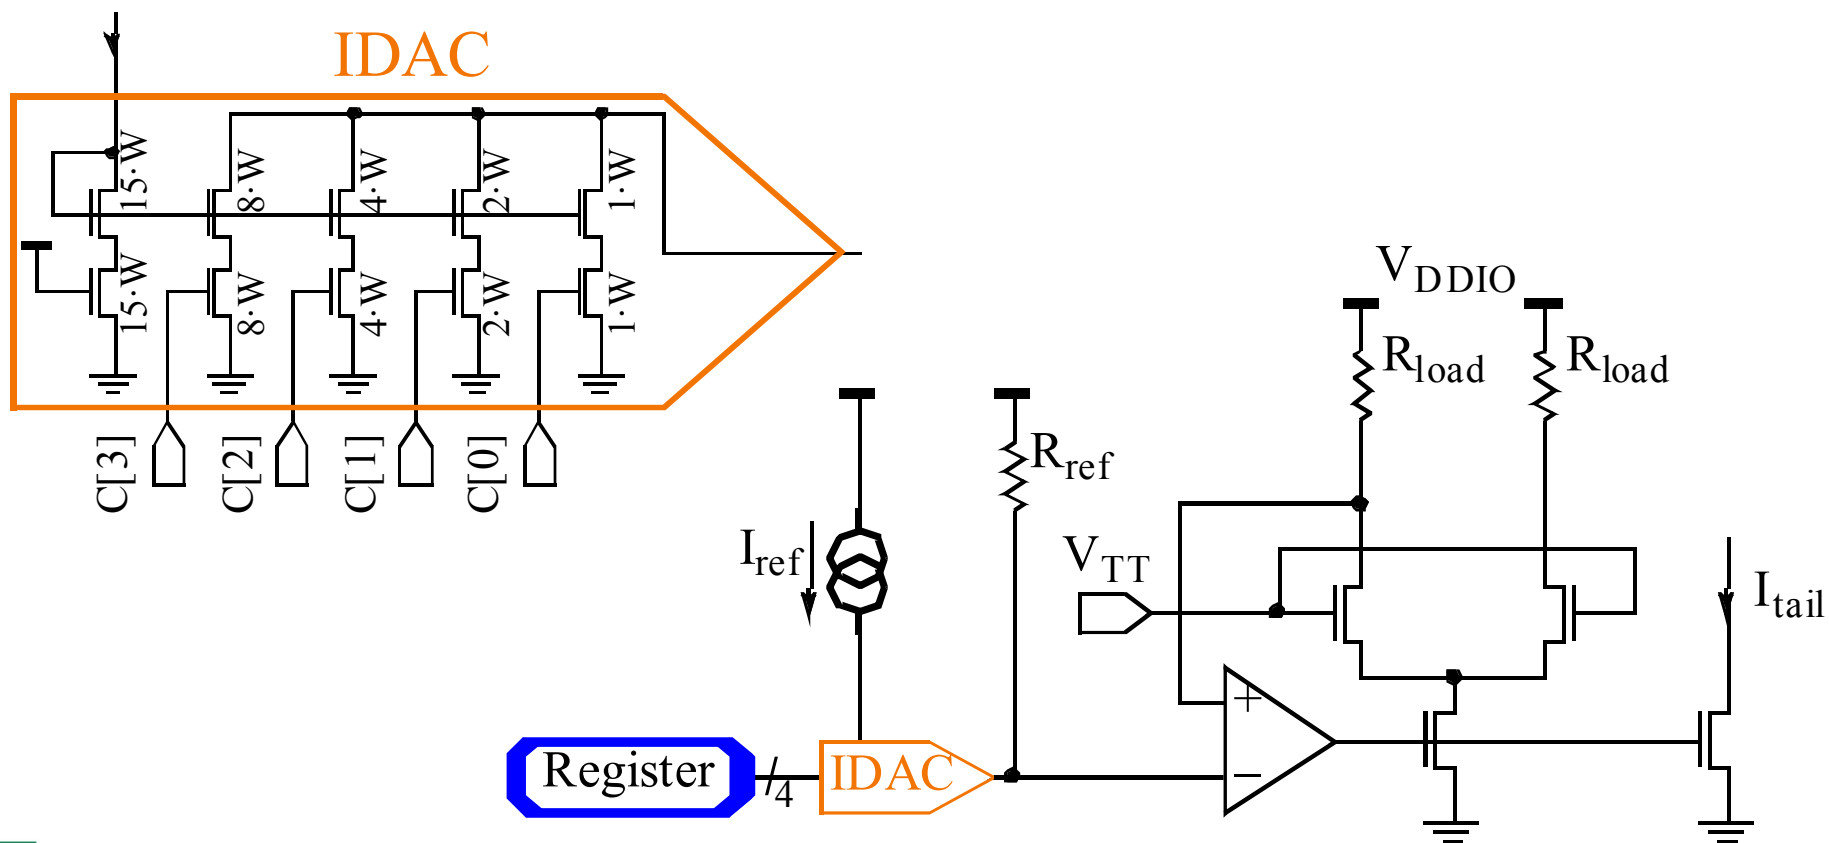
# DS DDR: Registers





# DS DDR: Servo Loop

- Use feedback to deal with PVT (servo)
- Use DACs to set the servo reference



# Summary

- Monolithic, Quad x86 cores
- 65nm SOI CMOS, 11 Cu layers
- Integrated Northbridge
- Enhanced 128b FPU
- 2MB shared L3
- 2 independent DDR channels, DDR2 & 3
- Enhanced HT links for 2 to 8Way MP
- Comprehensive thermal monitoring