

A Dual-Core Multi-Threaded Xeon[®] Processor with 16MB L3 Cache

Stefan Rusu
Intel Corporation
Santa Clara, CA




Outline

- Processor Highlights
- 65nm Process Technology
- Block Diagram
- L3 Cache Overview
- Sleep and Shut-off Modes
- Long-Le Transistor Usage
- Clock and Power Distribution
- Package Details
- DFT/DFM Features
- Thermal sensors
- Frequency Shmoo
- Summary





Processor Highlights



- Dual core, four threads
- 1 MB unified L2 per core
- 16 MB unified L3
- 435 mm² die size
- 1.328 B transistors
- 121 M transistors per core
- 3.4 GHz at 1.25 V and 150 W TDP
- 800 MT/s 3-load front-side bus interface
- Plugs in existing platforms

Largest cache and transistor count for an x86 processor

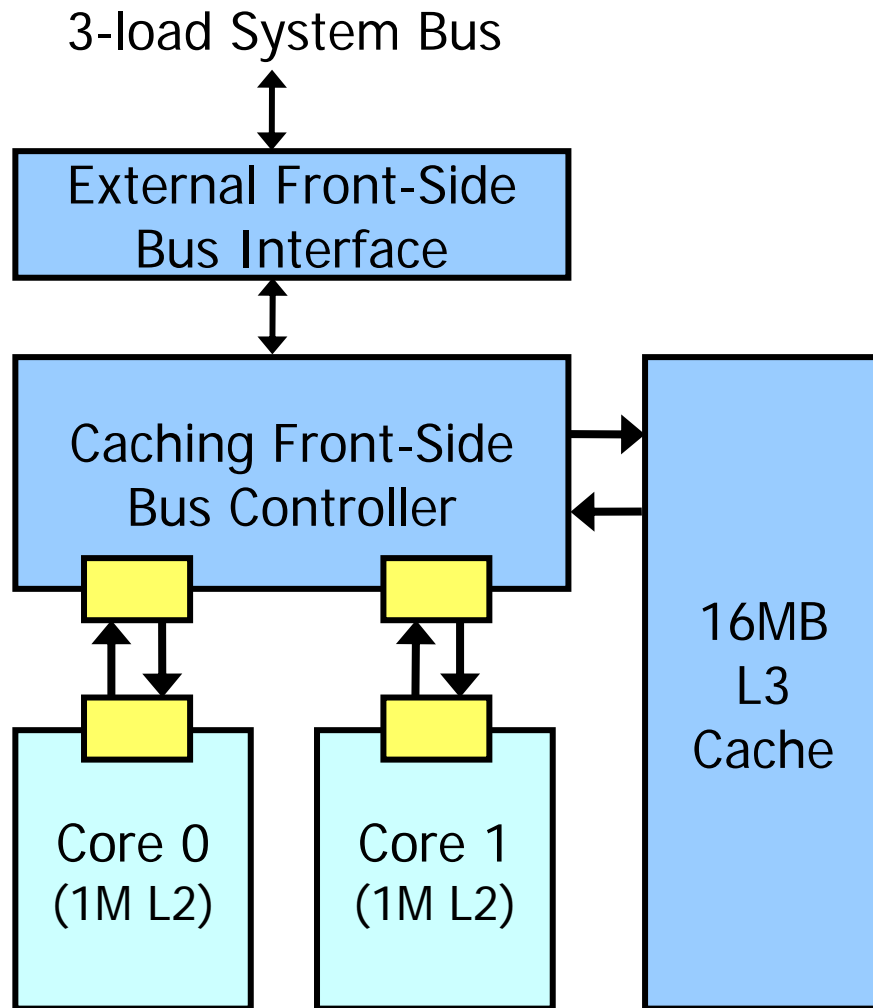
65nm Process Technology



- 1.2 nm gate oxide
- NiSi for low resistance
- Second generation strained silicon

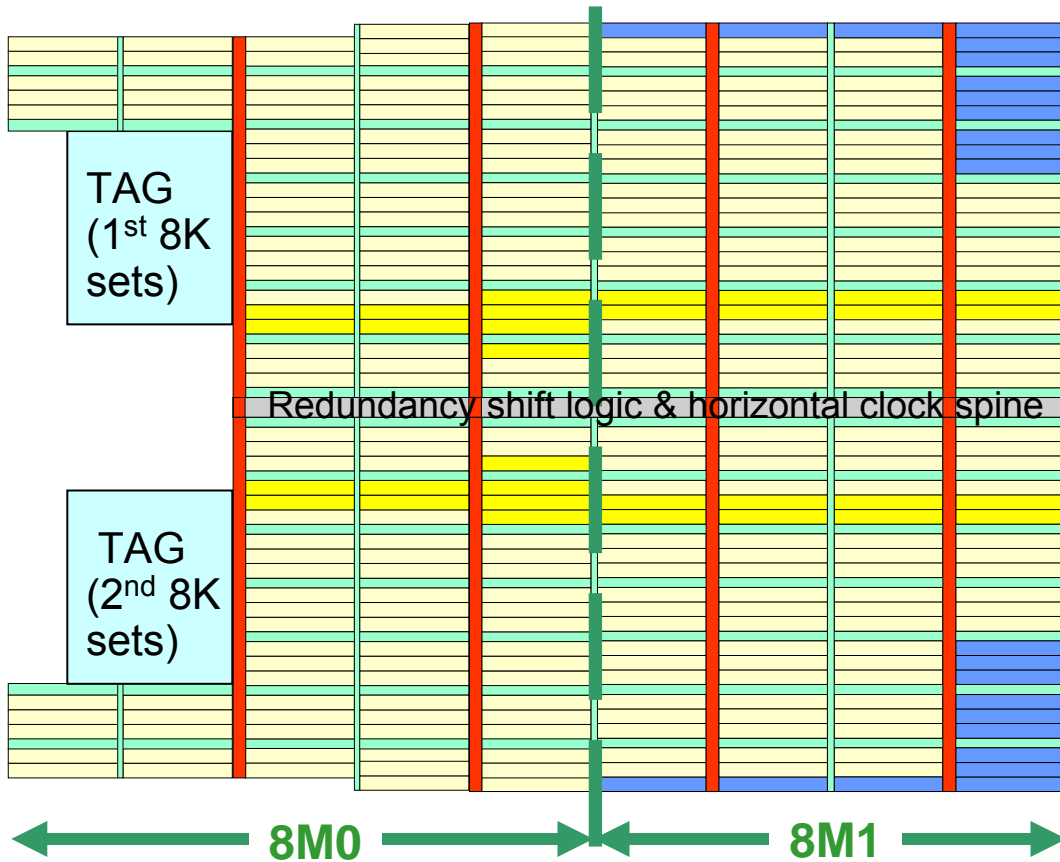
- 8 Cu interconnect layers with low-K CDO dielectric

Block Diagram




- Shared 16MB L3 cache
 - Better efficiency - one core can use more than half of the total cache
 - No need for coherency traffic between caches
- Support for Intel Virtualization Technology
- Hyper-Threading Technology
- Enhanced Intel SpeedStep® Technology

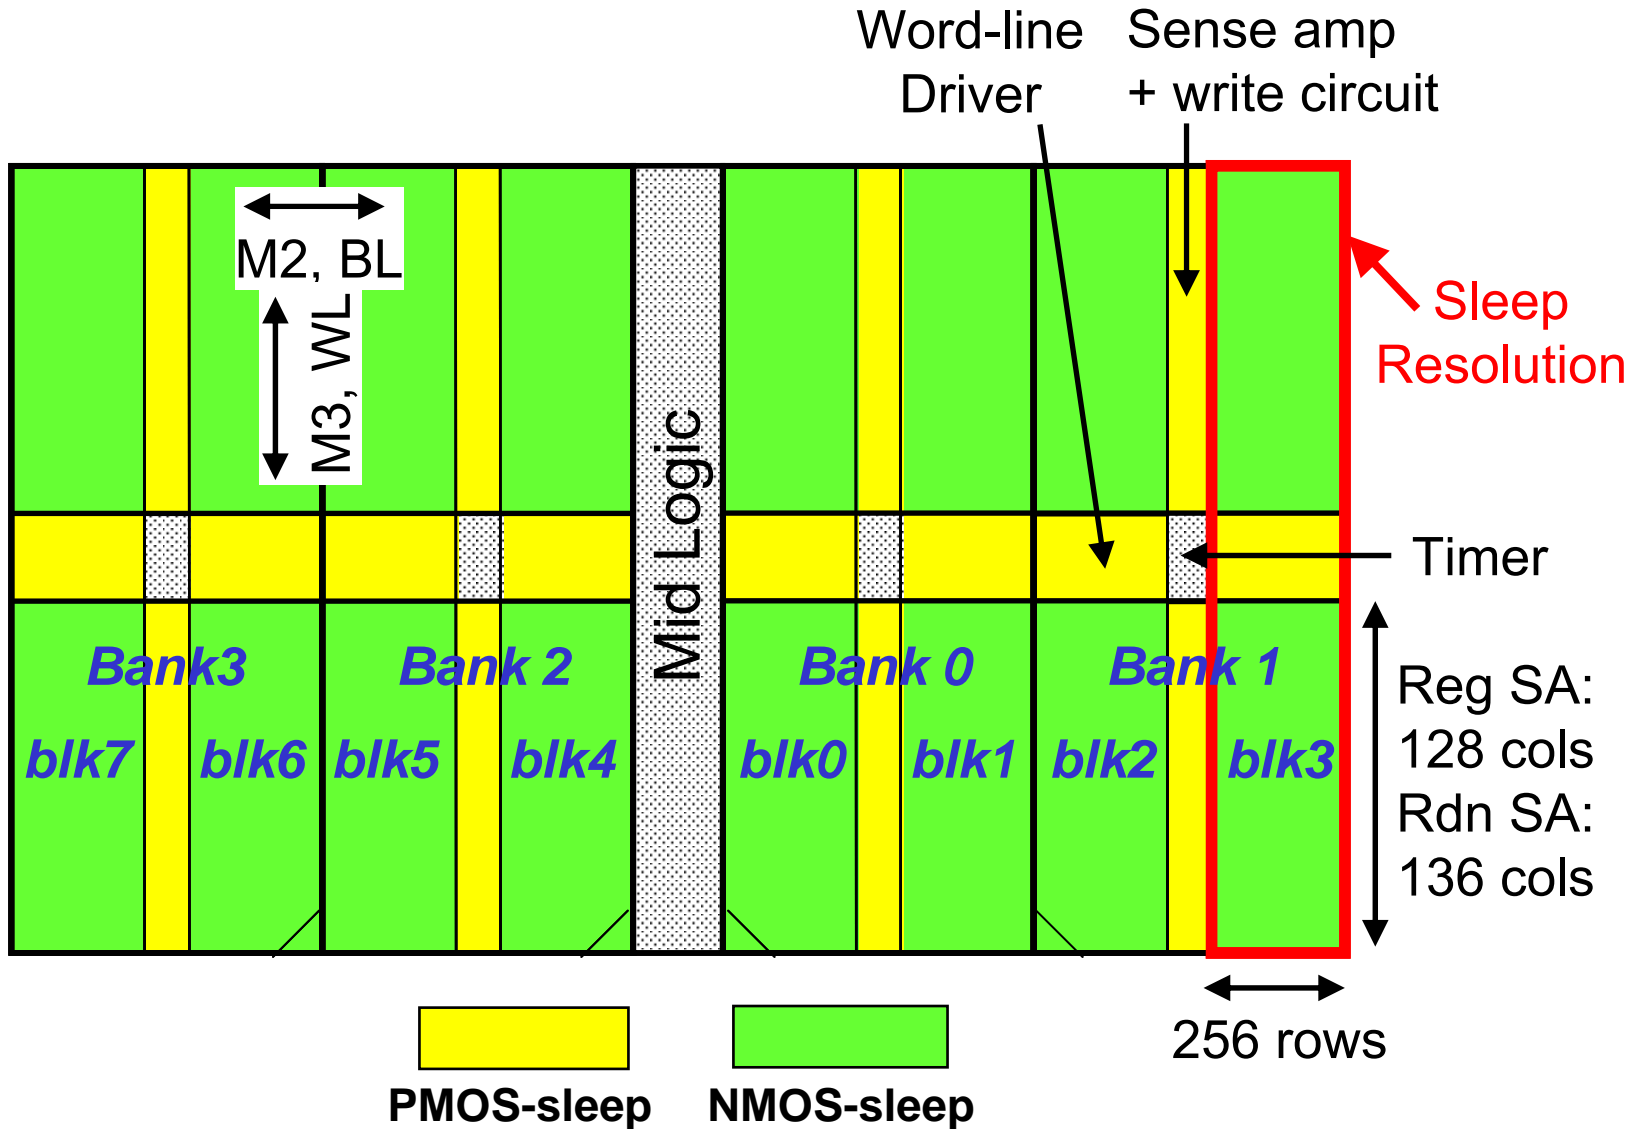
L3 Cache Floorplan



- 0.624 μm^2 bit cell
- 0.75 Watts / MB average power
- Only 0.8% of all array blocks powered-up for each access
- 256 regular 64kB sub-arrays
- 32 redundancy 68kB sub-arrays



L3 Sub-Array and Sleep Partitioning



Column Redundancy Features


- Data Array

- Two redundancy columns for each 290-bit chunk
- Repair up to two random defects in each cache line



- Tag Array

- One redundancy column for each 36-bit tag line
- Repair one random defect in each entry


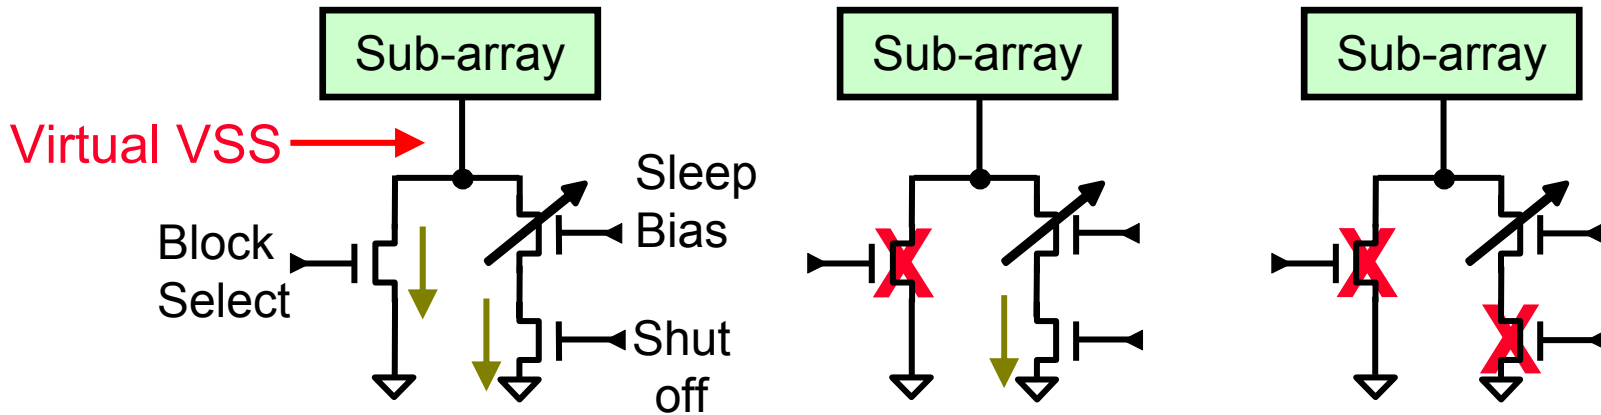


L3 Cache Sleep and Shut-off Modes

Active Mode

Sleep Mode


Shut-off Mode



Leakage Shut-off Infrared Images

16MB SKU


All 16MB in
sleep mode



8MB SKU


8MB in
sleep mode

8MB in
shut-off mode




Shut-off feature reduces the leakage of the
8MB disabled sub-arrays by about 3W

Leakage Mitigation: Long-Le Transistors



- All transistors can be either nominal or long-Le
- Most library cells are available in both flavors
- Long-Le transistors are about 10% slower, but have 3x lower leakage
- All paths with timing slack use long-Le transistors

Long-Le Transistors Usage Map




Long-Le Transistors Summary

Percentage of Long-Le device width excluding RAM arrays:

Cores


Nominal
46%



Long-Le
54%

Uncore


Nominal
24%



Long-Le
76%

Moore's Law will continue to double transistors every 2 years
New Trend ► To reduce sub-threshold leakage, most devices will be slower and only a handful of transistors will be fast

Clock Domains



Legend:



Core



PLL



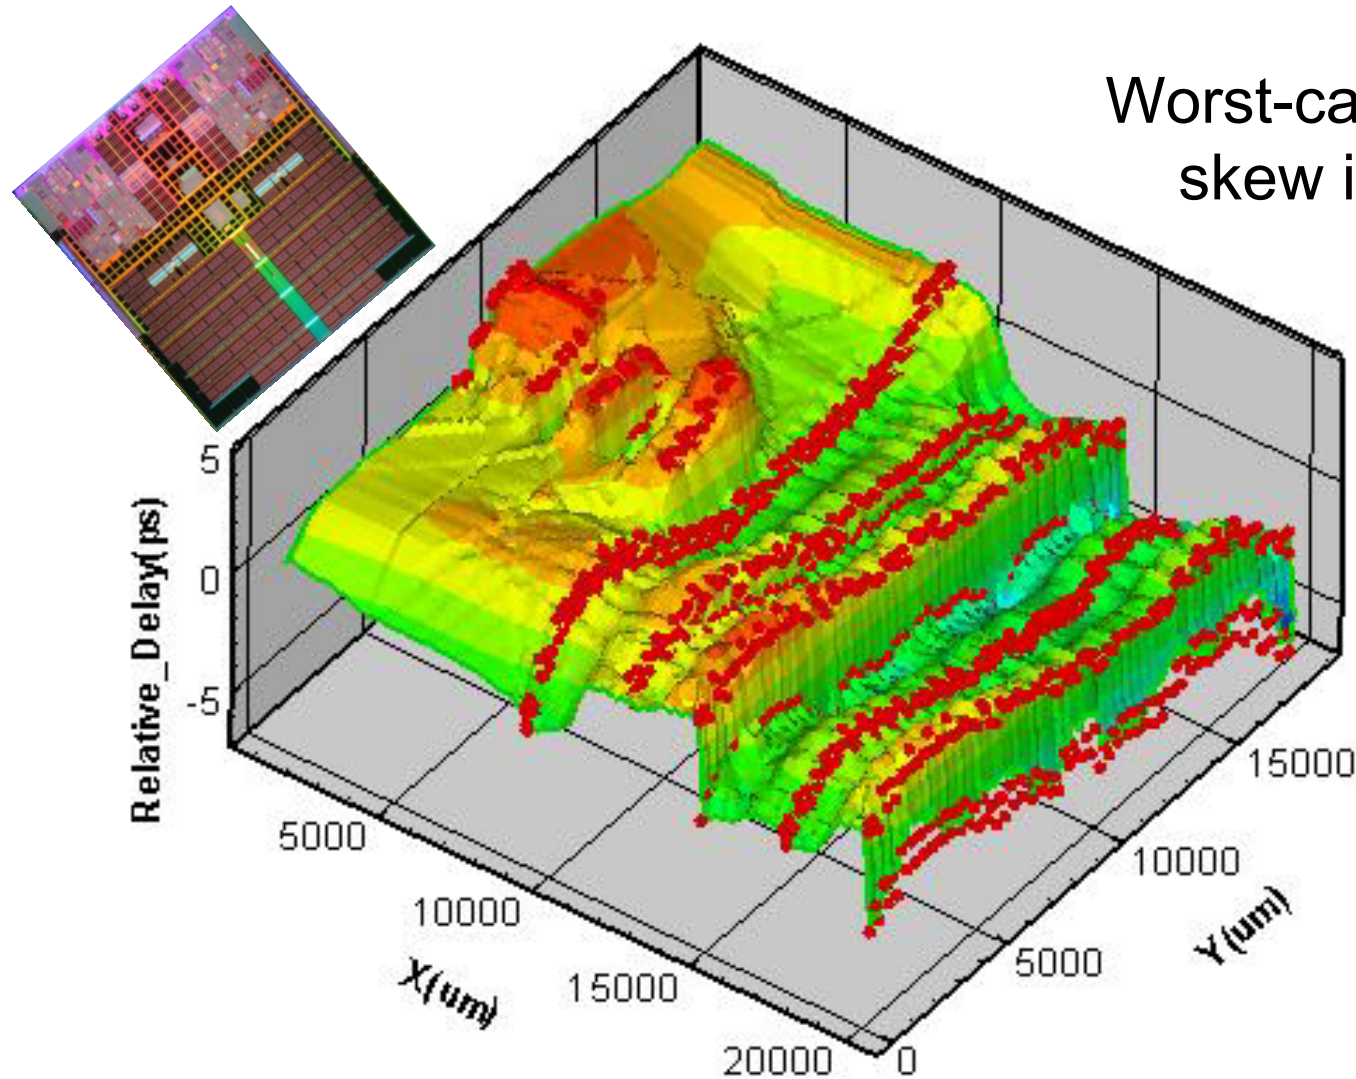
Uncore



I/O




Global Clock Skew Profile



Worst-case global skew is 11ps

More details on the clock distribution in paper 21.2

Voltage Domains



Legend:



Core



PLL



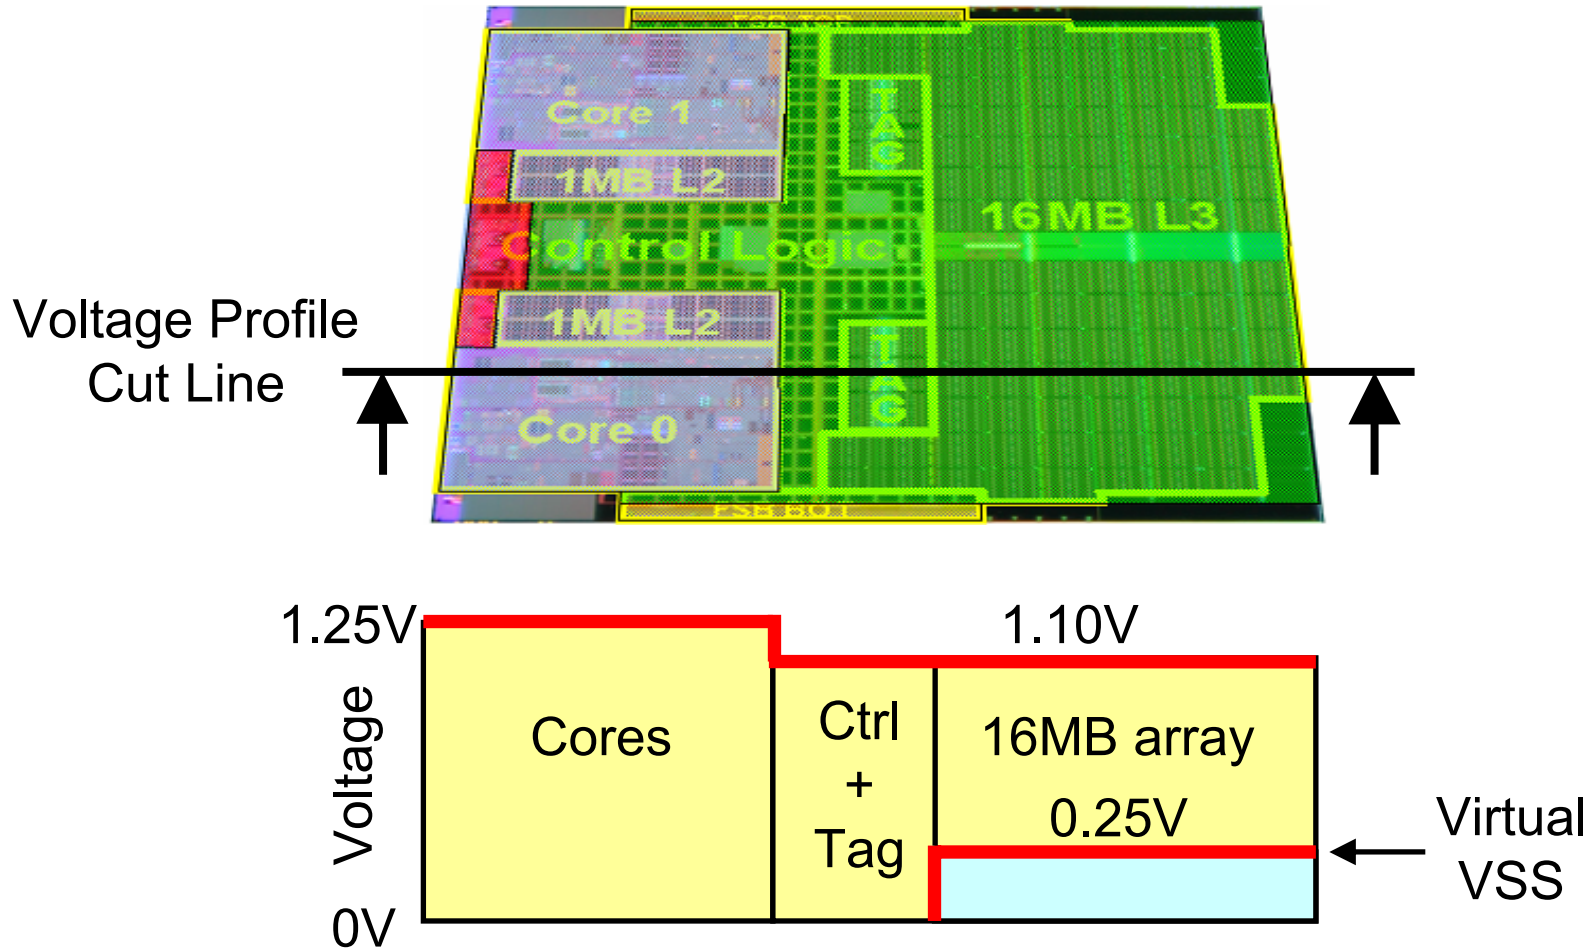
Uncore



I/O



Voltage Profile

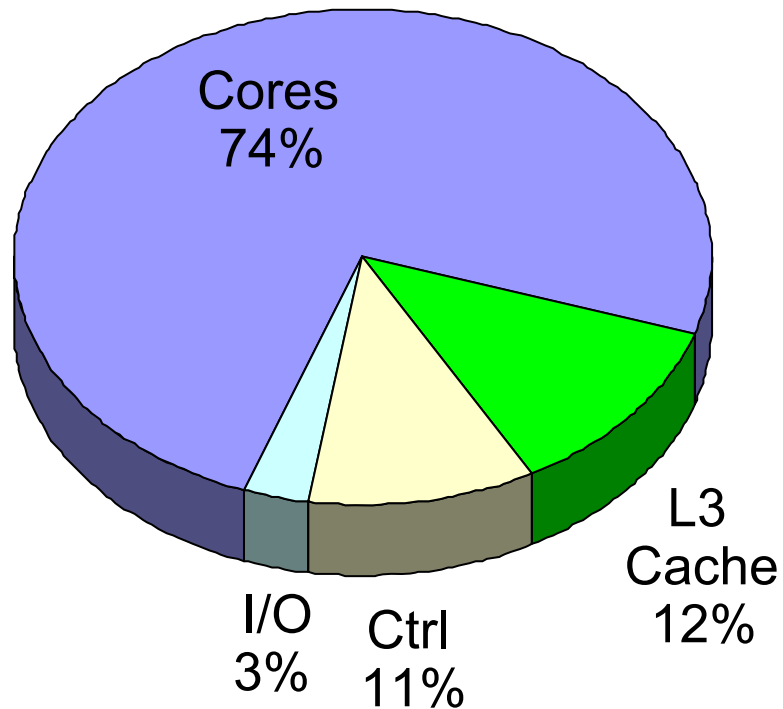


Cache sleep function enables separate voltage knob

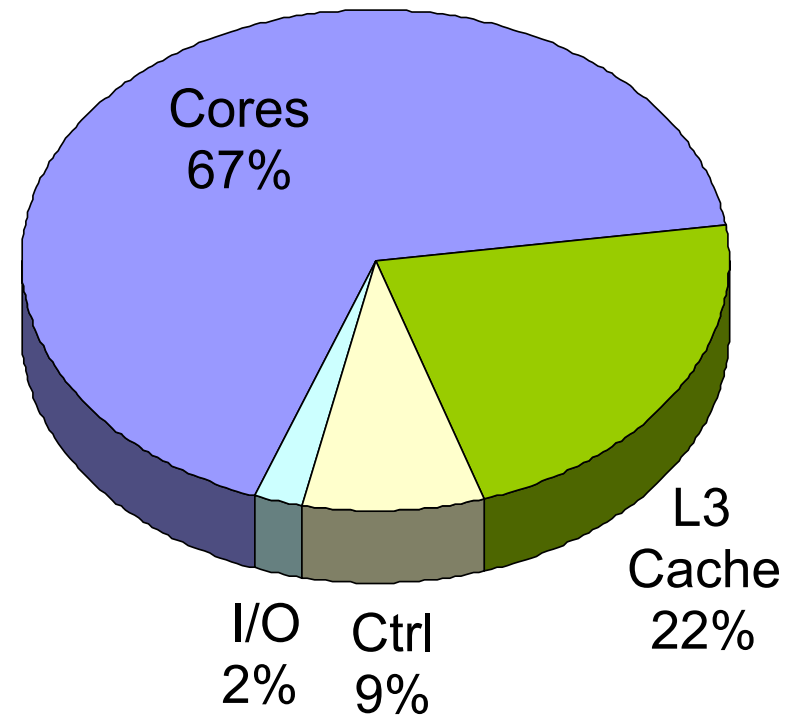


Power and Leakage Breakdown

Total Power Breakdown

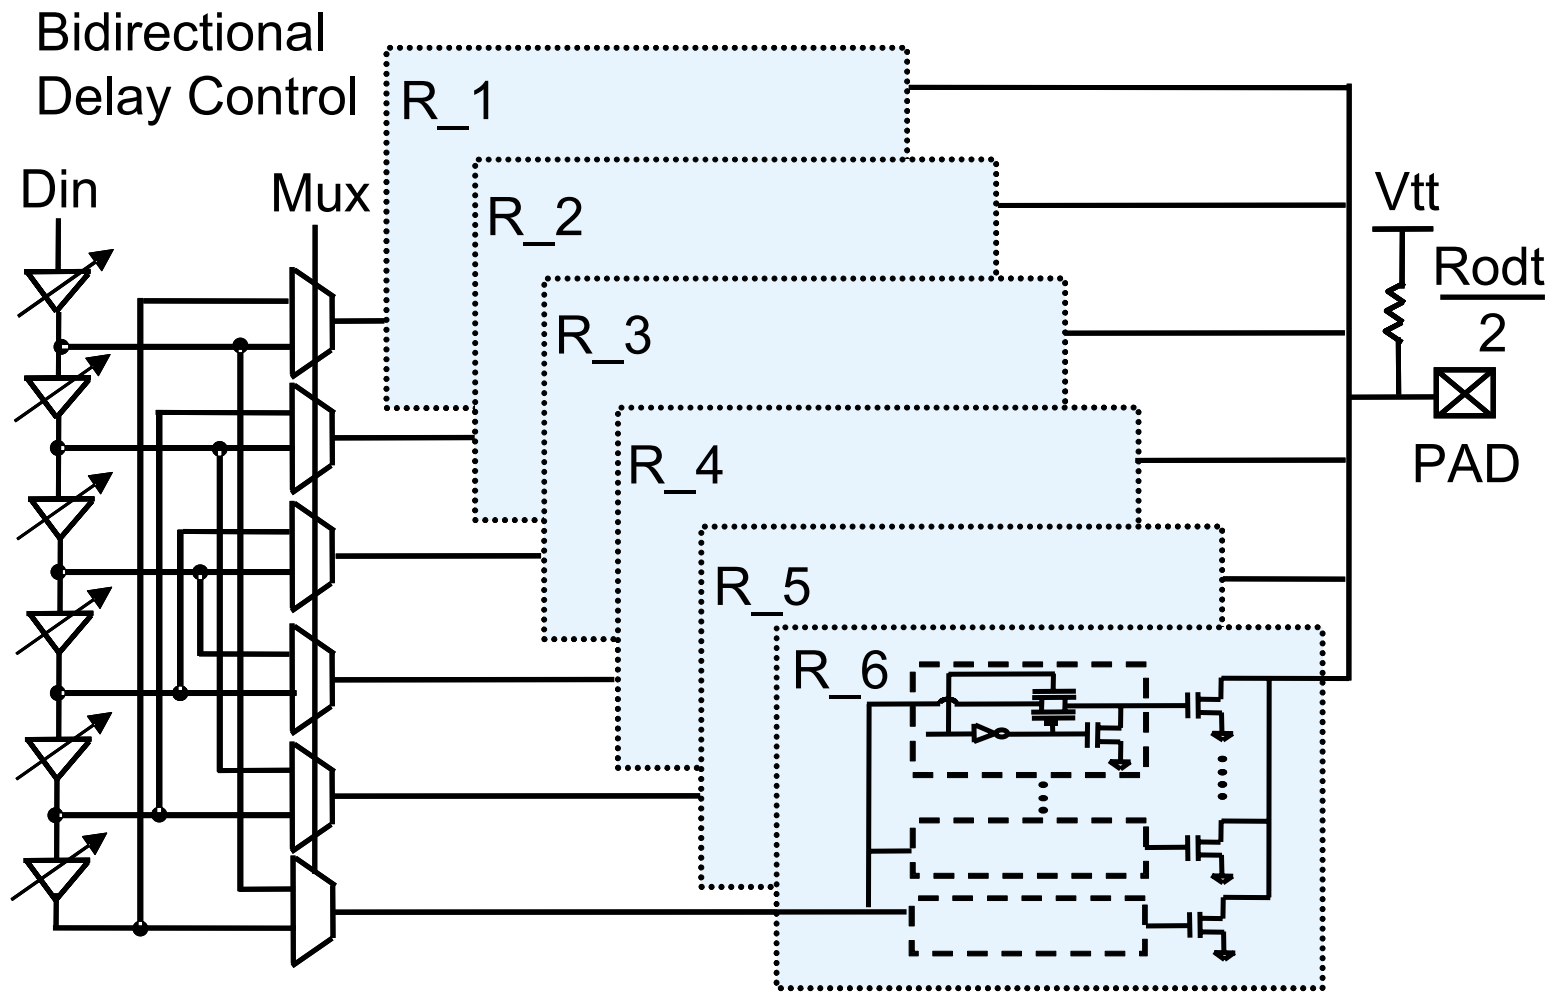


Leakage Breakdown




Leakage accounts for about 30% of the total power

Symmetric I/O Pre-driver Circuit




C4 Bump Map



- 13164 C4 bumps
- Perfectly uniform bump pitch over the entire die (including I/O buffers) to improve epoxy underfill

Package Details



- 12 layers organic substrate (53.3 mm/side)
- 4-4-4 stacking
- Integrated heat spreader (38.5 mm/side)
- 604 total pins
- 366 signal I/Os
- System management components on package

Design for Test and Debug Features

Die-level DFT/DFM

- Parallel structural core test with XOR
- Scan and observability registers (scan-out)
- Three TAP controllers (core0, core1, uncore)
- Within-die process monitors
- On-die clock shrink

L3 cache DFT/DFM

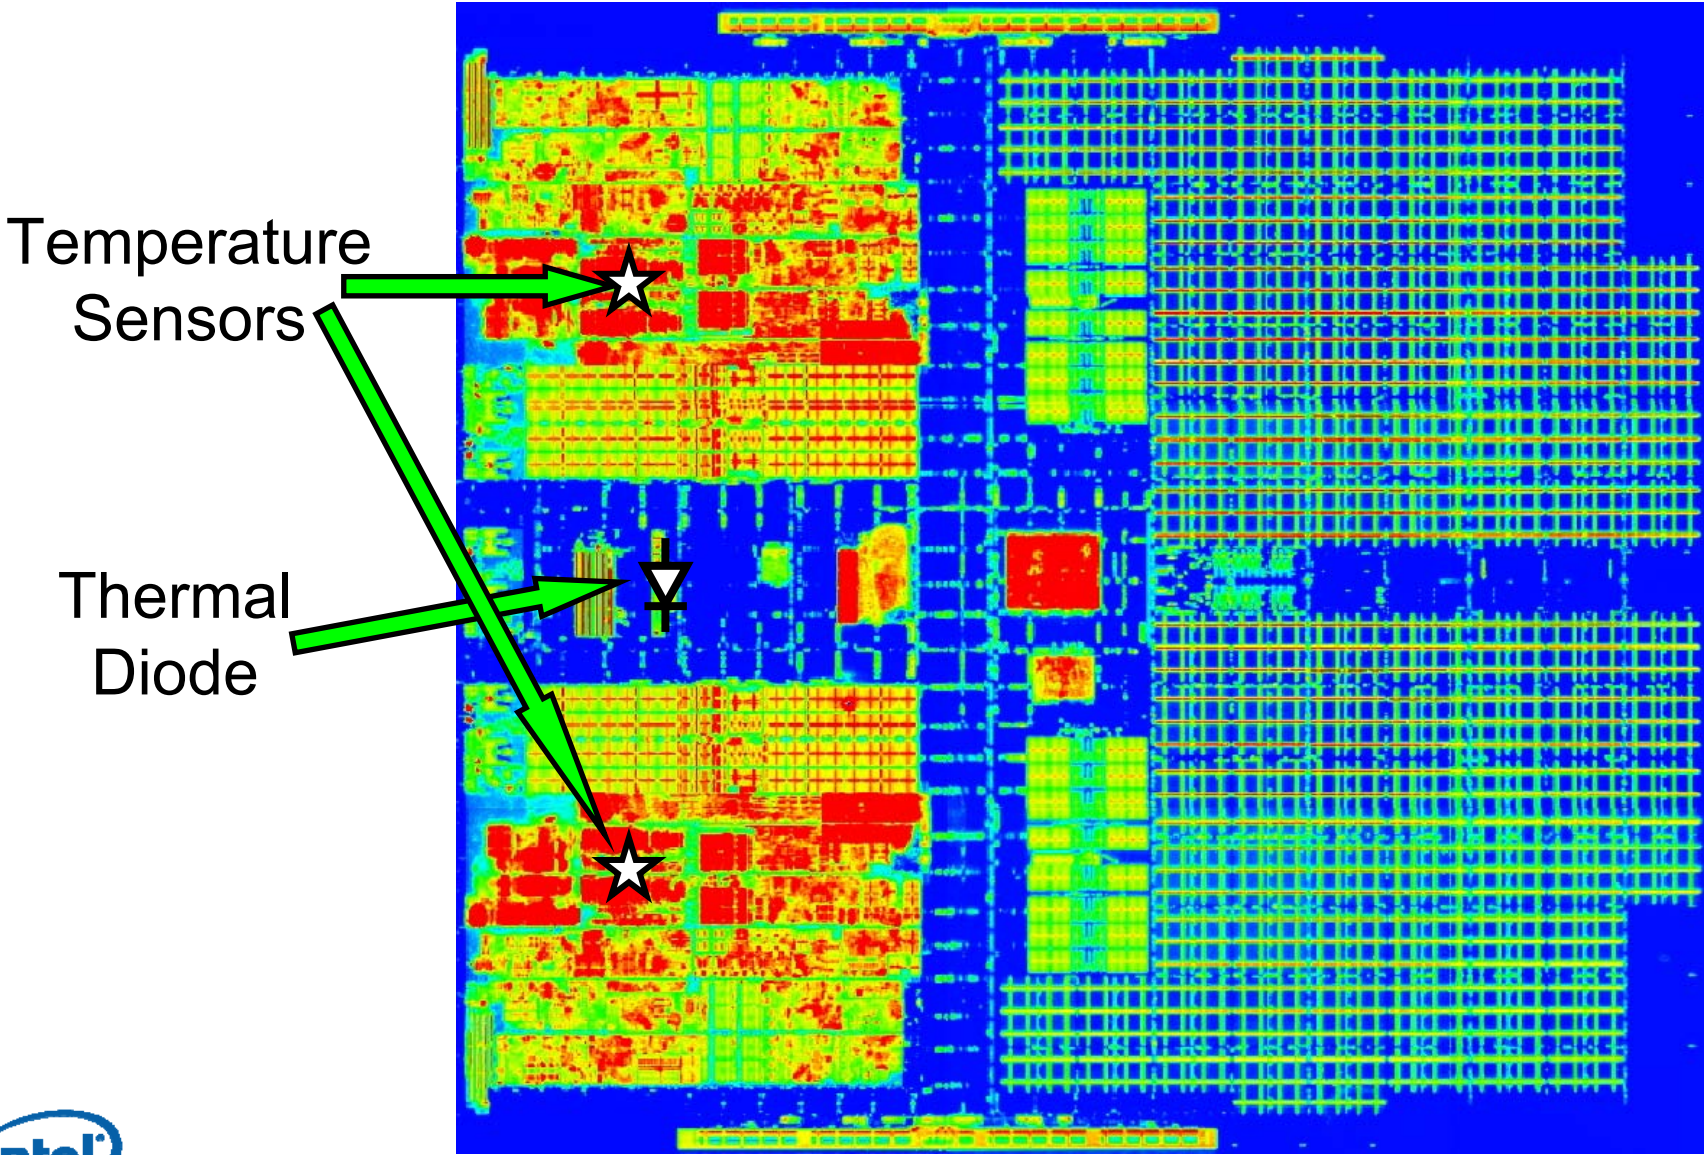
- Built-in pattern generator (PBIST)
- Programmable weak-write test
- Low-yield analysis
- Stability test mode
- 32-entry cache line disable (Pellston)

FSB DFT/DFM

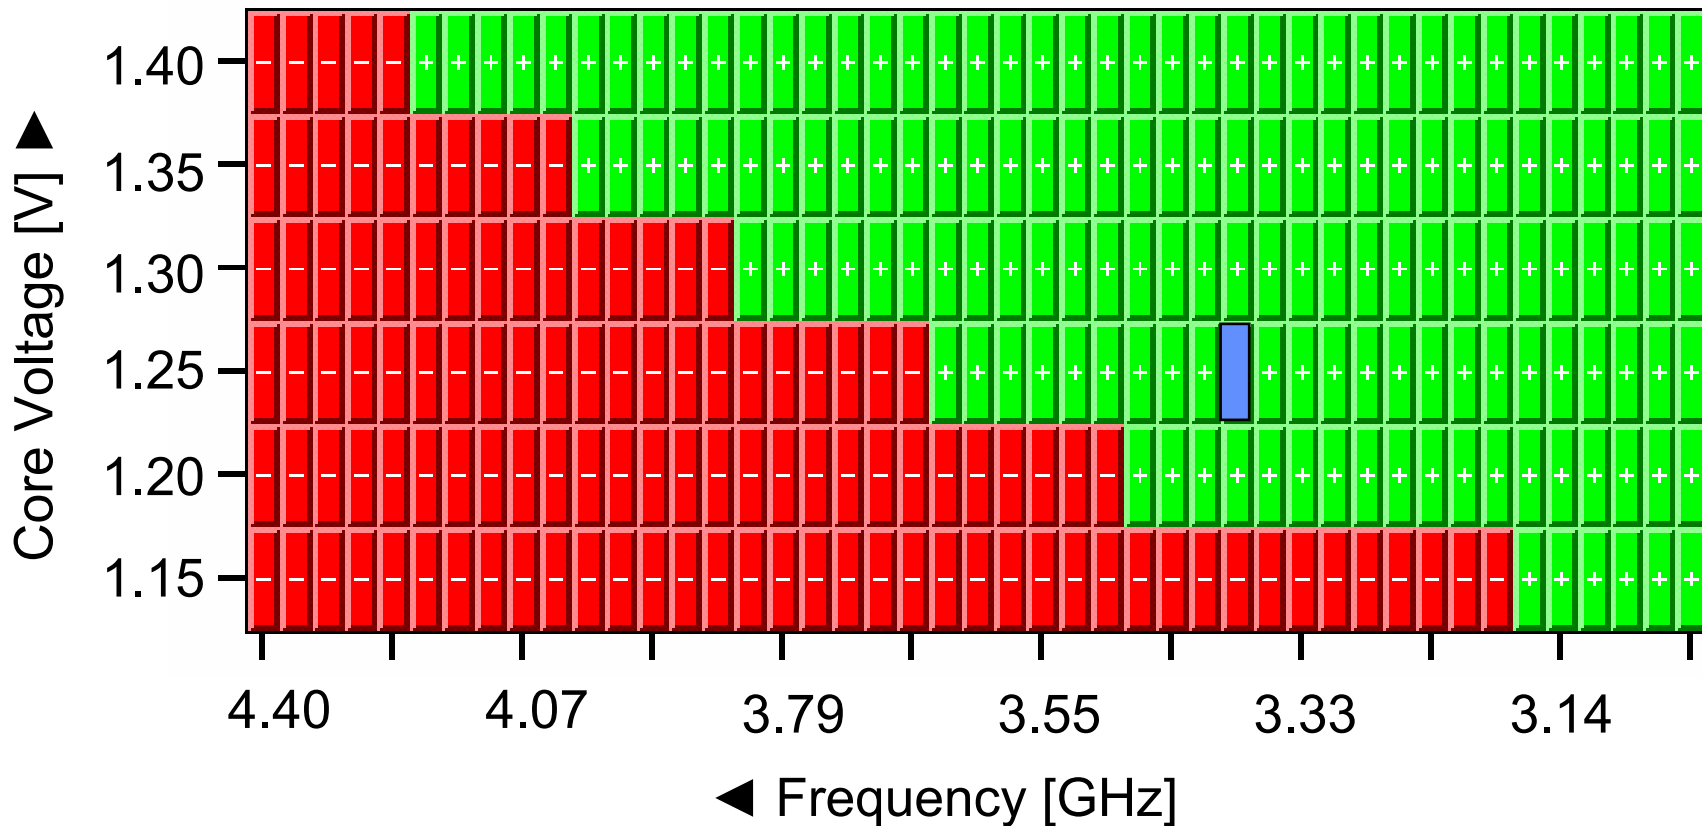
- I/O loopback
- I/O test generator



Thermal Sensors



Frequency Shmoo



Target operating point 3.4GHz at 1.25V core voltage and 150W TDP envelope

Summary

- 65nm dual-core, four-thread Xeon[®] Processor with on-die 16MB L3 cache
 - Shared L3 cache ▶ best fit for server processors
 - Virtualization and Hyper-Threading Technology
- Leakage reduction circuit techniques
 - Massive Long-Le usage
 - N and P sleep transistors in L3 cache
 - L3 leakage shut-off mode ▶ saves power in lower cache size SKUs
- Multiple voltage and clock domains to reduce active power and leakage

