



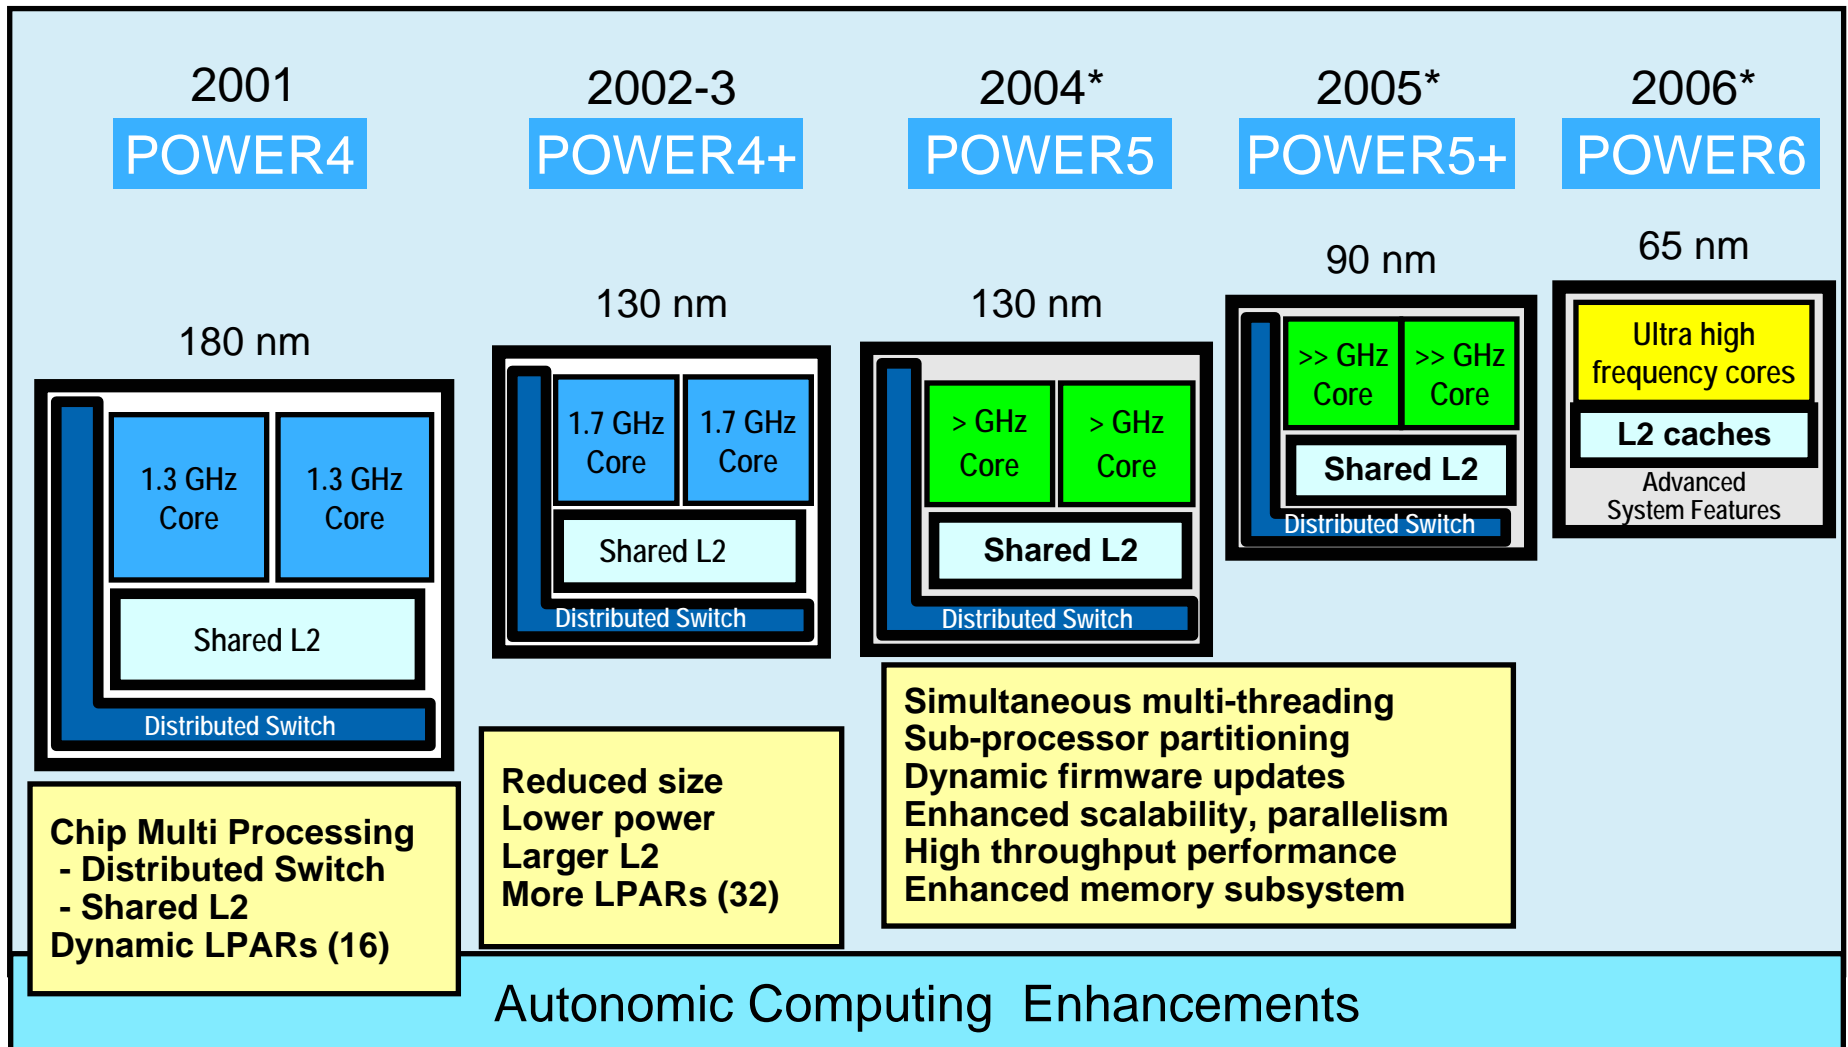
IBM's POWER5 Micro Processor Design and Methodology

Ron Kalla
IBM Systems Group

Outline

- POWER5 Overview
- Design Process
- Power


POWER Server Roadmap



*Planned to be offered by IBM. All statements about IBM's future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only.


POWER5

- Technology: 130nm lithography, Cu, SOI
- 389mm² 276M Transistors
- Dual processor core
- 8-way superscalar
- Simultaneous multithreaded (SMT) core
 - ▶ Up to 2 virtual processors per real processor




Multi-threading Evolution


Single Thread



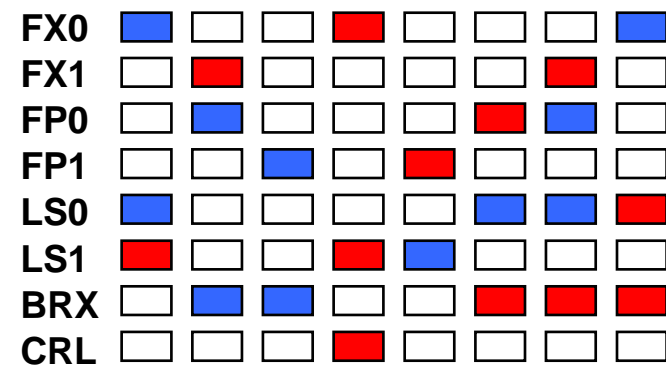
Coarse Grain Threading



Fine Grain Threading




Simultaneous Multi-Threading



█ Thread 0 Executing
 █ Thread 1 Executing
 □ No Thread Executing

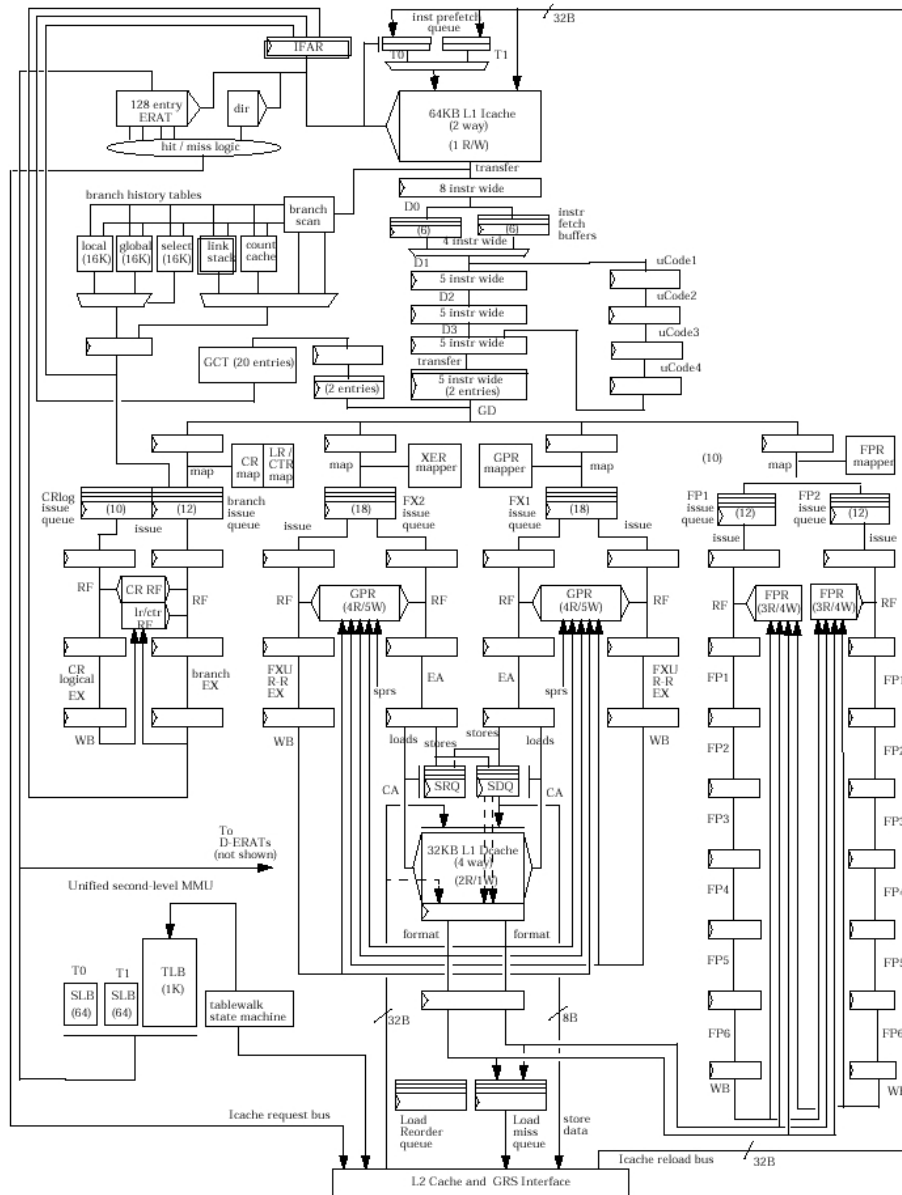
Thread Priority

- Instances when unbalanced execution desirable
 - ▶ No work for opposite thread
 - ▶ Thread waiting on lock
 - ▶ Software determined non uniform balance
 - ▶ Power management
 - ▶ ...
- Solution: Control instruction decode rate
 - ▶ Software/hardware controls 8 priority levels for each thread

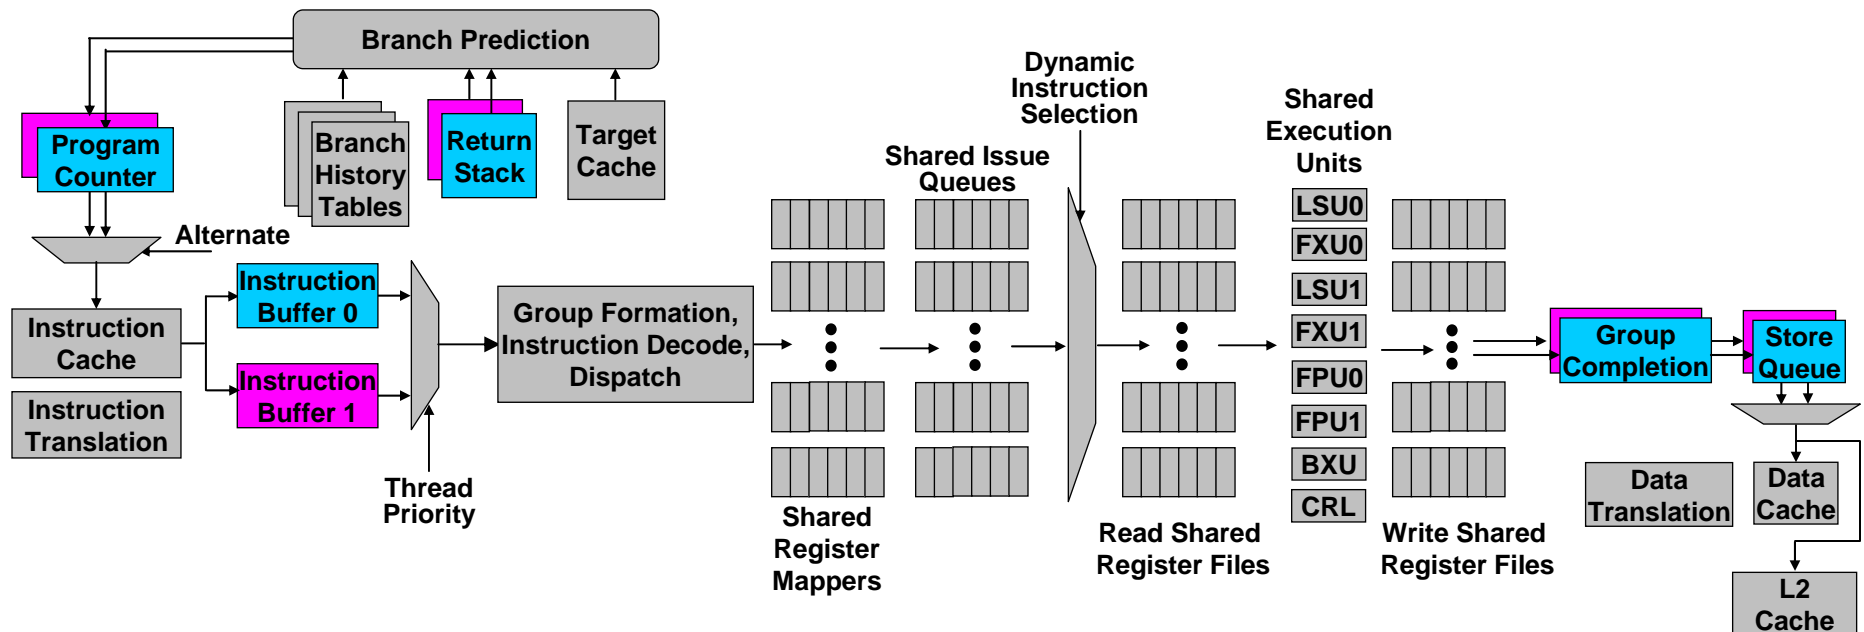
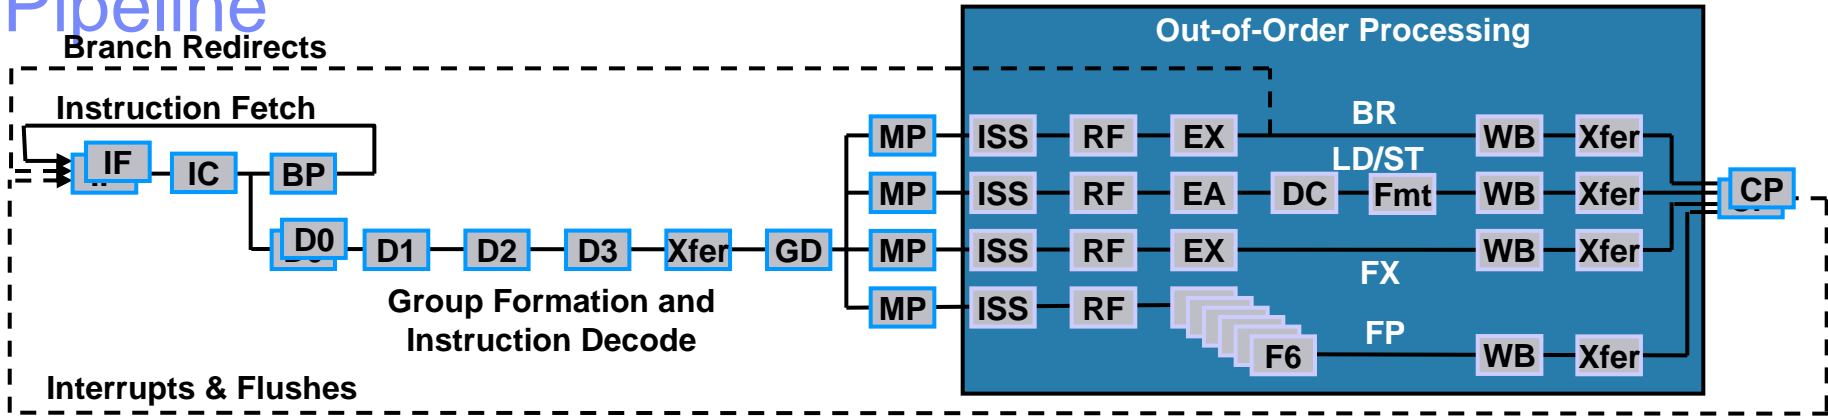


Terminology

- PowerPC Addresses
 - ▶ Effective>(SLB)>Virtual>(Page Table)>Real>(LPAR)>Physical
- Instruction Execution
 - ▶ I-fetch
 - ▶ Decode
 - ▶ Dispatch
 - ▶ Issue
 - ▶ Finish
 - ▶ Complete




Multithreaded Instruction Flow in Processor Pipeline



Shared by two threads
 Resource used by thread 0
 Resource used by thread 1

Resource Sizes

- Analysis done to optimize every micro-architectural resource size
 - ▶ GPR/FPR rename pool size
 - ▶ I-fetch buffers
 - ▶ Reservation Station
 - ▶ SLB/TLB/ERAT
 - ▶ I-cache/D-cache
- Many Workloads examined
- Associativity also examined




Results based on simulation of an online transaction processing application

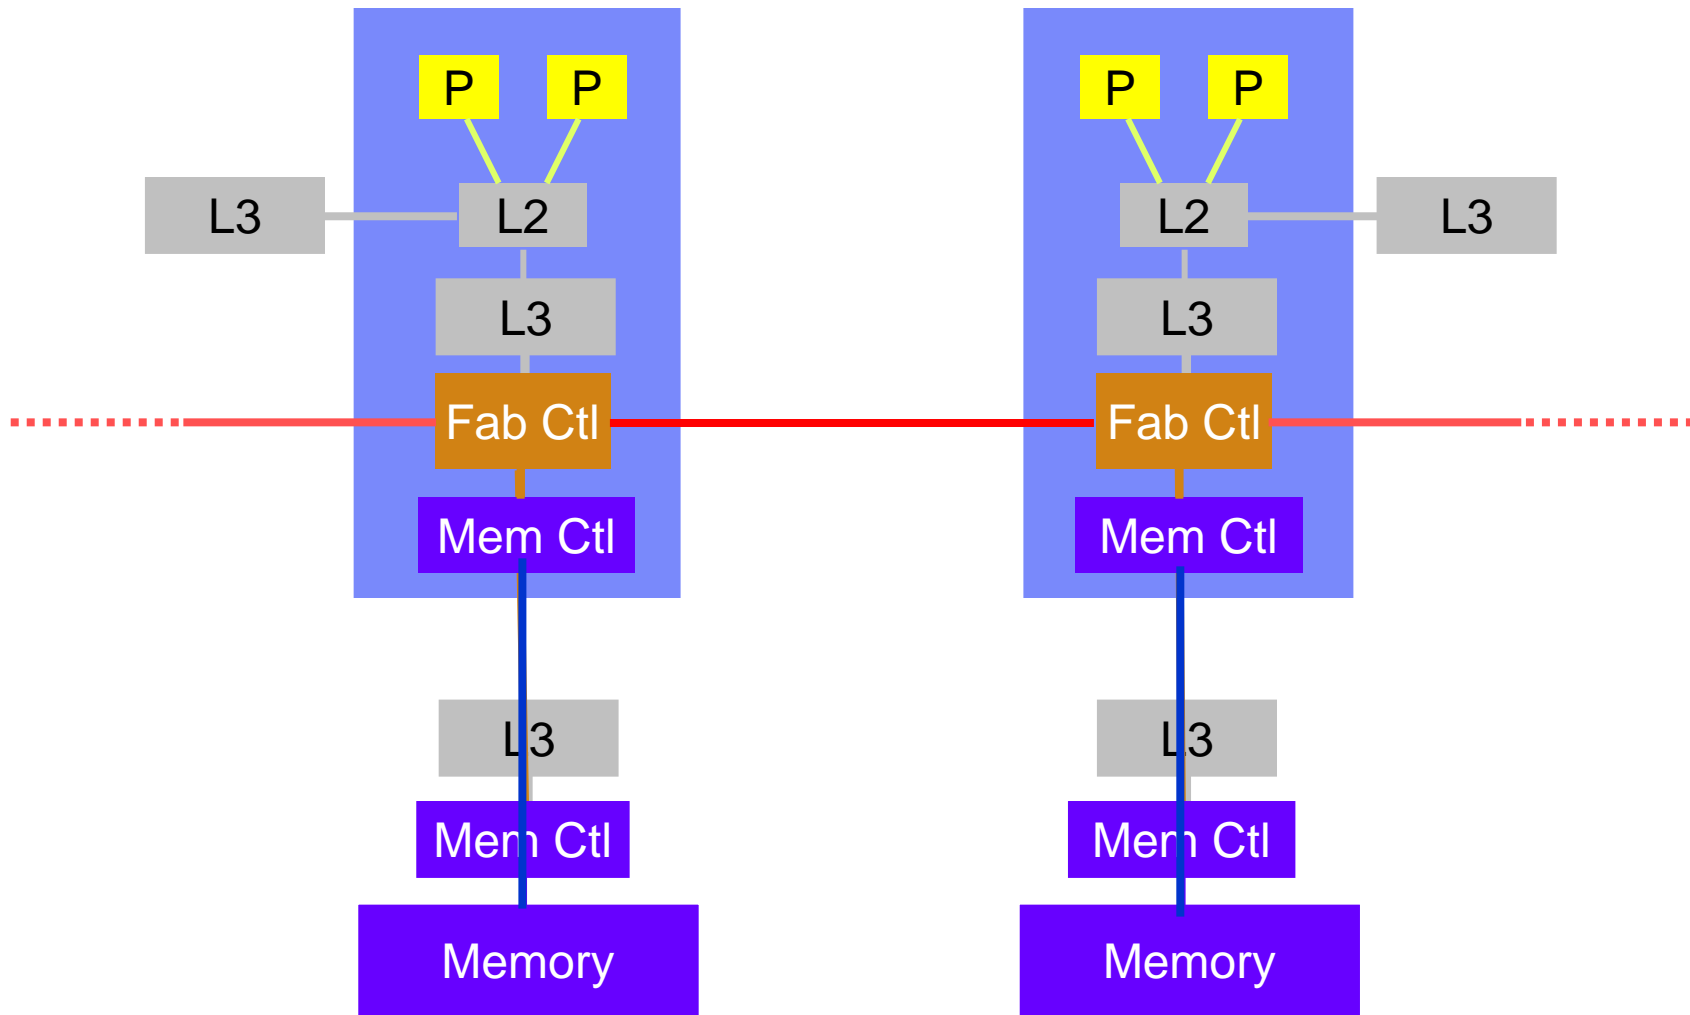
Vertical axis does not originate at 0

Single Thread Operation


- Advantageous for execution unit limited applications
 - ▶ Floating or fixed point intensive workloads
- Execution unit limited applications provide minimal performance leverage for SMT
 - ▶ Extra resources necessary for SMT provide higher performance benefit when dedicated to single thread
- Determined dynamically on a per processor basis



Modifications to POWER4 System Structure




16-way Building Block

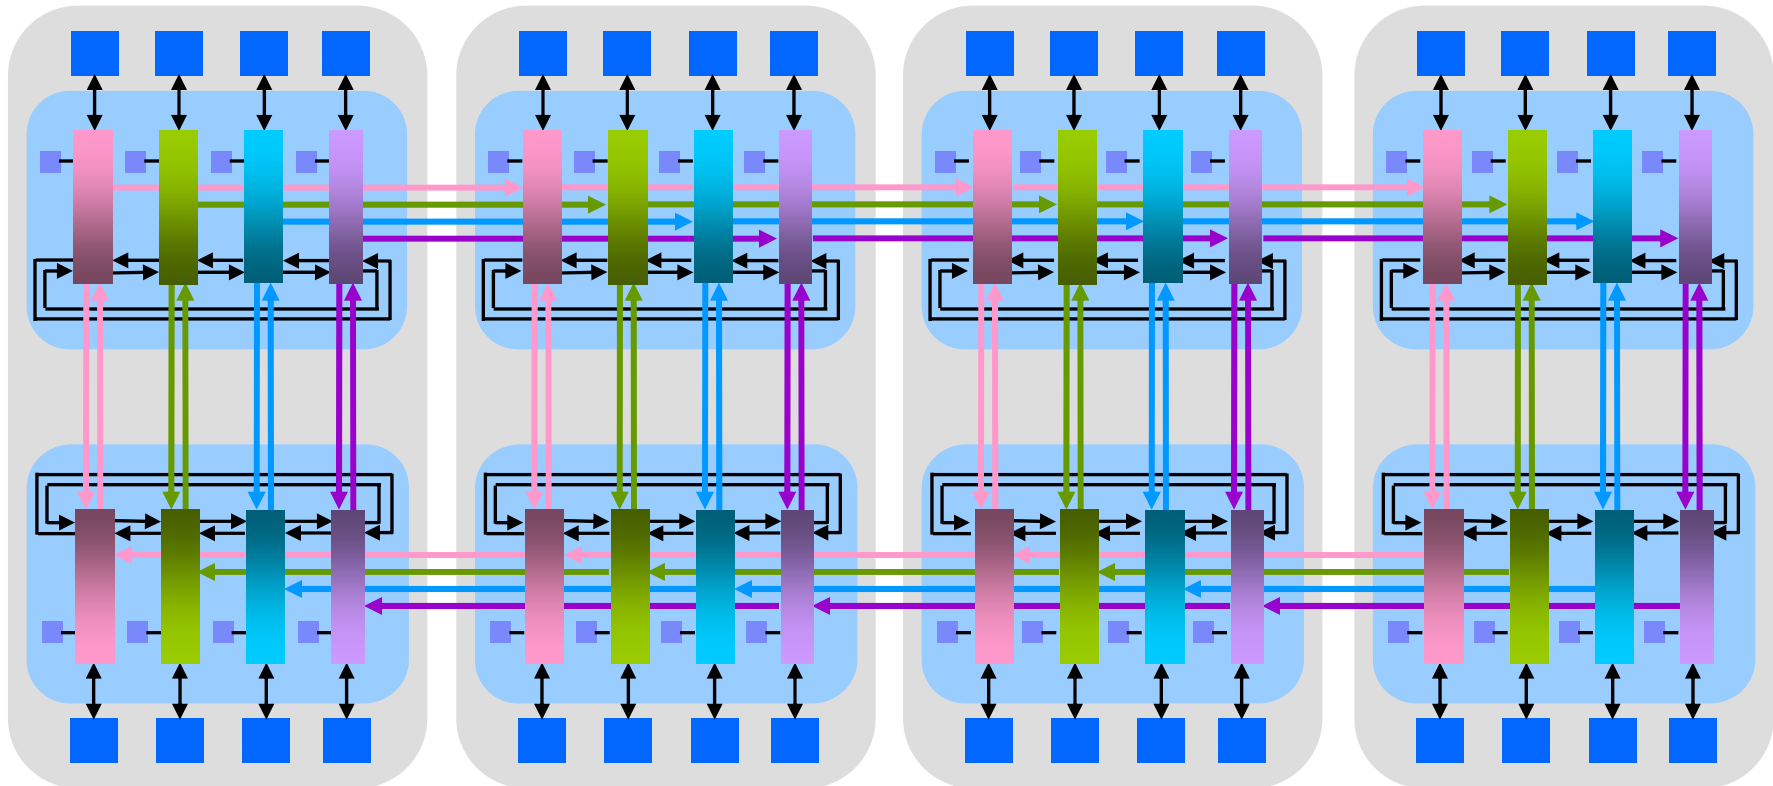


POWER5 Multi-chip Module

- 95mm × 95mm
- Four POWER5 chips
- Four cache chips
- 4,491 signal I/Os
- 89 layers of metal



64-way SMP Interconnection



Interconnection exploits *enhanced distributed switch*

- All chip interconnections operate at half processor frequency and scale with processor frequency

Design Process

- Concept Phase (~10 People/4 months)
 - ▶ Competitive assessment
 - ▶ Customer requirements
 - ▶ Technology assessment
 - Try to match technology introduction with products
 - ▶ AREA/POWER estimates
 - Cost
 - ▶ Schedule
 - ▶ Staffing
 - ▶ Lots or executive/technical reviews

Design Process (cont.)

- High Level Design Phase (~50 People/6 months)
 - ▶ Micro-architecture
 - ▶ Cycle Time closed
 - Cross section of critical paths
 - Power closed
 - ▶ Area/Power budgets established
 - ▶ Arrays architected
 - ▶ All Unit I/O defined timing contract in place
 - ▶ Performance model meets concept phase targets
 - ▶ HLD Exit review.
 - ▶ Everything is now in place to effectively engage full team

Design Process (cont.)

- Implementation Phase (~200 People/12-18 months)
 - ▶ Write VHDL
 - ▶ Schematic entry for circuits
 - ▶ Simulation
 - Block>unit>core>chip>system
 - ▶ Integration process begins
 - RLM and Custom macros > unit > core > chip
 - ▶ Design rule checking, (can't trust human's)
 - ▶ Bulk simulation
 - ▶ Check recheck and check again
- RIT (Release Information Tape)

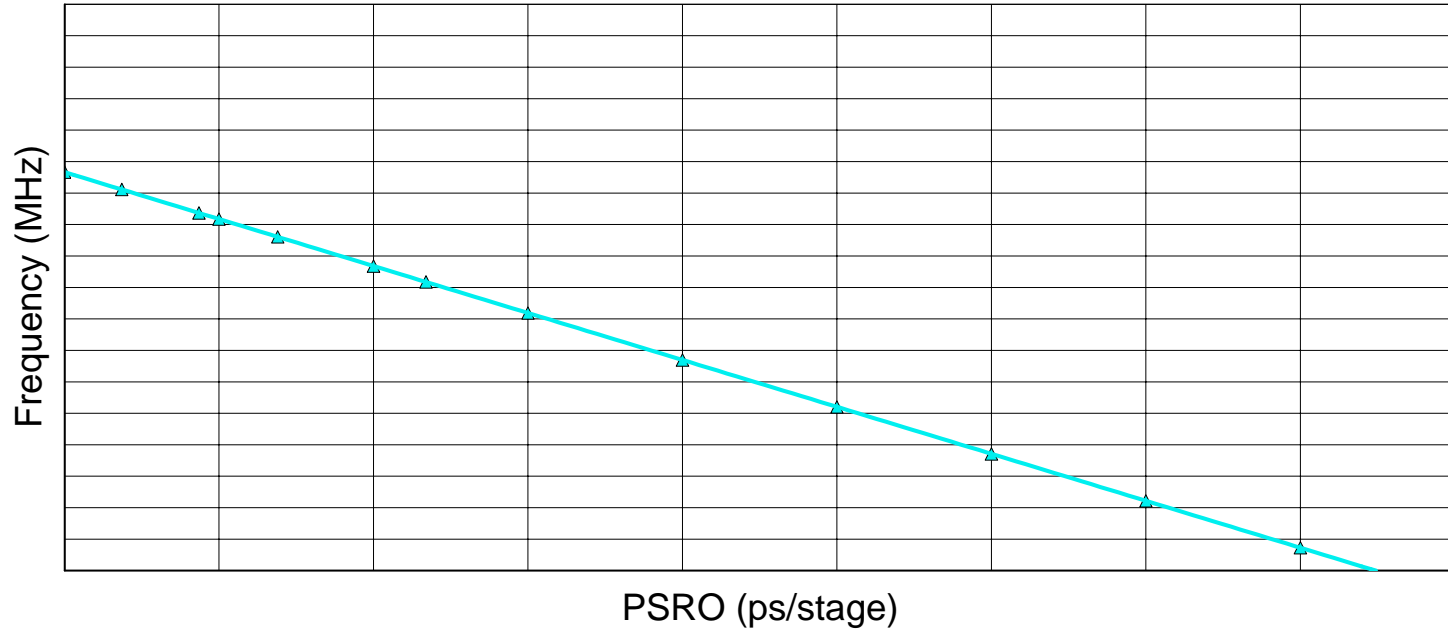
Bring-up 100's People 12-18 Months

- Starts at Wafer test.
- AVPs
- Random test programs
- Boot operating system
 - ▶ OS based testing
- Formal testing (new groups of people)
 - ▶ POWER (Will discuss later)
 - ▶ Performance testing
- Additional RITs if necessary
- Release to mfg
- GA (Hurray)
- Field Support

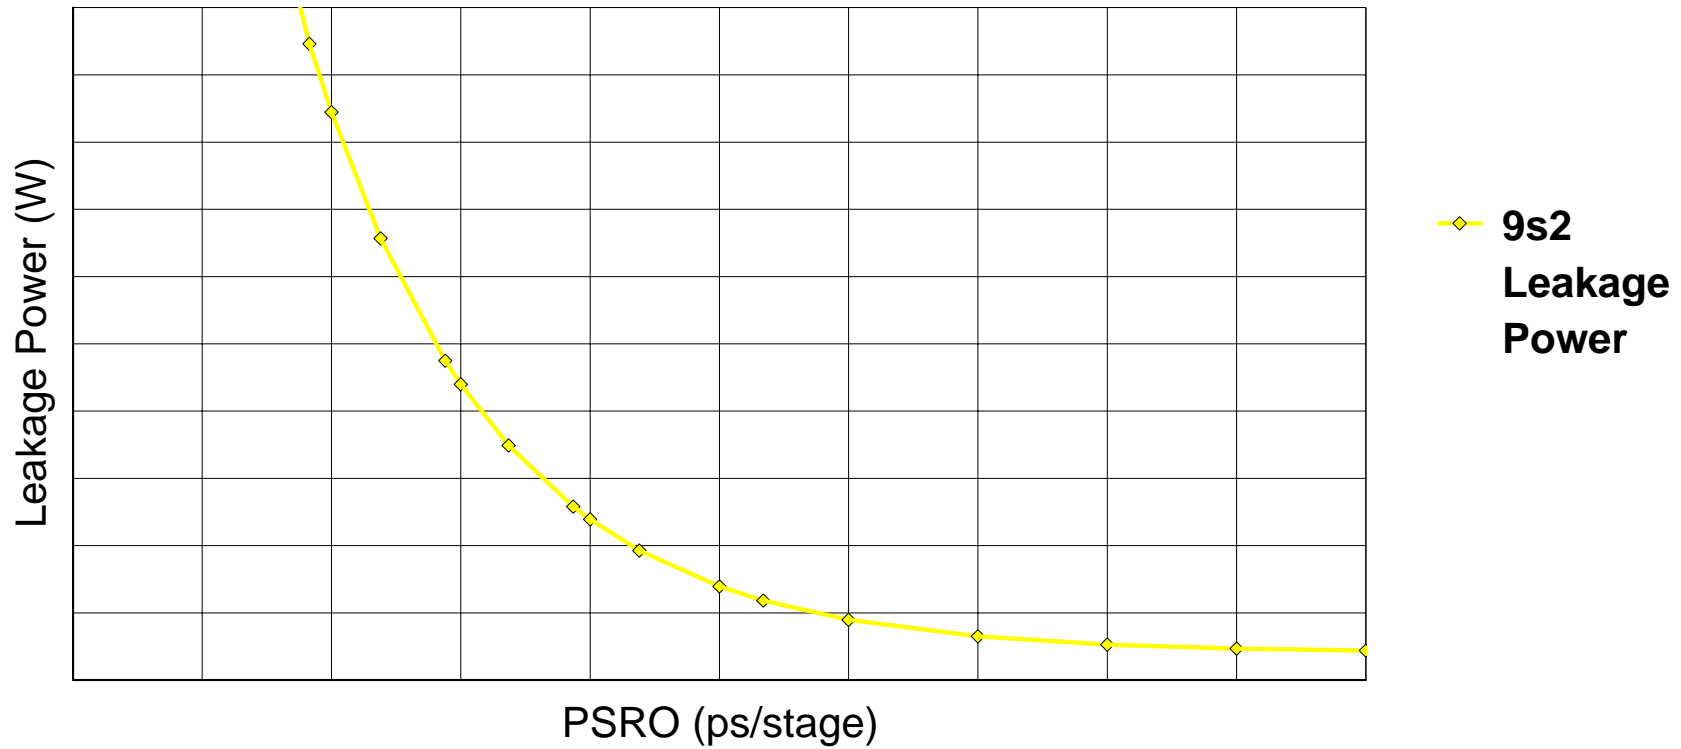
Power

- Micro-Processor Power Complex equation.
 - ▶ DC Power (IDDQ)
 - Voltage
 - Temperature
 - Speed of Chip (PSRO) Processor speed Ring Oscillator.
 - ▶ AC Power
 - Function of CV^2f X Switching Factor.
 - Workload Dependent.


GR DD2.0 Fmax vs. 1.2v, 25C PSRO Unguardbanded, 1.3v, 85C Fmax




Power 5 Leakage Power vs. PSRO 1.3 V and 85C



Power 5 AC Power vs. Frequency 1.3 V and 85 C



Chip Power vs Frequency 1.3V 85C



Other SMT Considerations

- Power Management
 - ▶ SMT Increases execution unit utilization
 - ▶ Dynamic power management does not impact performance
- Debug tools / Lab bring-up
 - ▶ Instruction tracing
 - ▶ Hang detection
 - ▶ Forward progress monitor
- Performance Monitoring
- Serviceability