PRESENTED BY
CAHNERS MICRODESIGN RESOURCES

Power4: A Dual-CPU Processor Chip

Jim Kahle
IBM

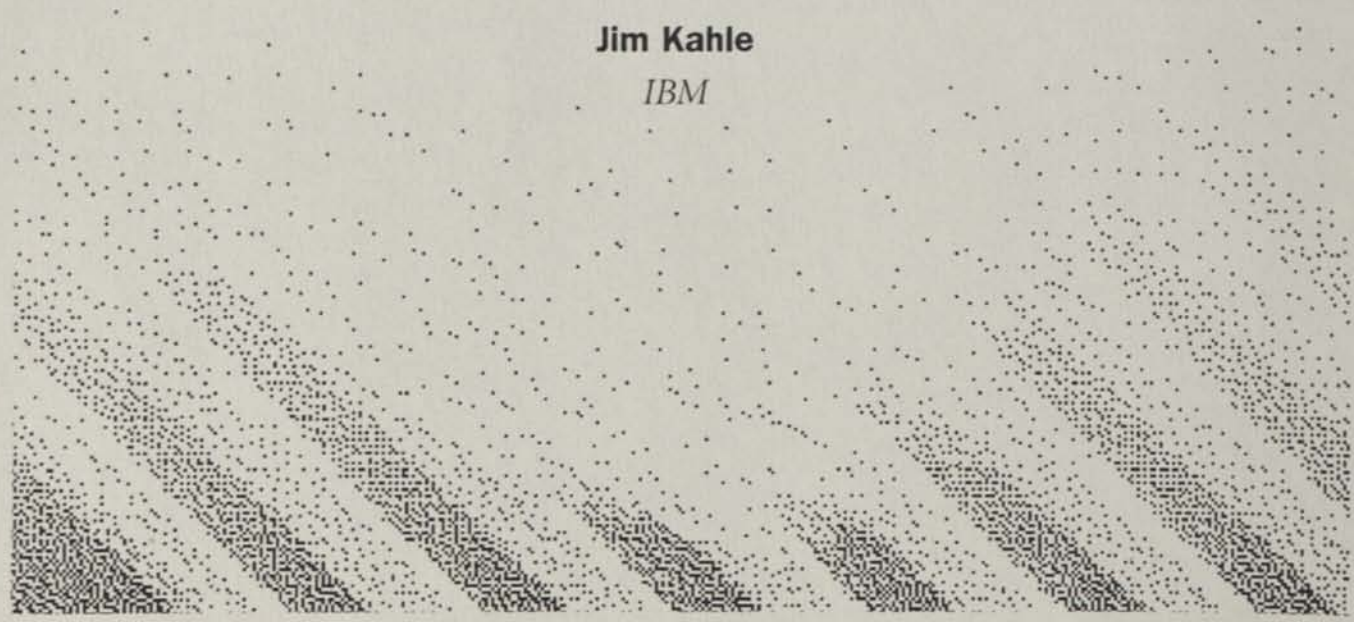

POWER4: A Dual-CPU Processor Chip

Jim Kahle
Chief Architect
Senior Technical Staff Member
IBM

IBM

Server Design Attributes

- **Large system parallelism and concurrent execution**
- **Balanced cache and memory subsystem**
- **Software design from the OS through the application**
- **Fully compatible and very scalable**
- **Highly available and serviceable**

IBM POWER4

Server Workload Characteristics

- **Commercial**
  - Large database footprints
  - Small record access
  - Random access patterns
  - Sharing / Thread communication
- **Technical**
  - Structured data
  - Large data movement
  - Predictable strides
  - Minimal data reuse

E-business applications include attributes from both commercial and technical workloads

IBM POWER4

Bandwidth Drives Performance

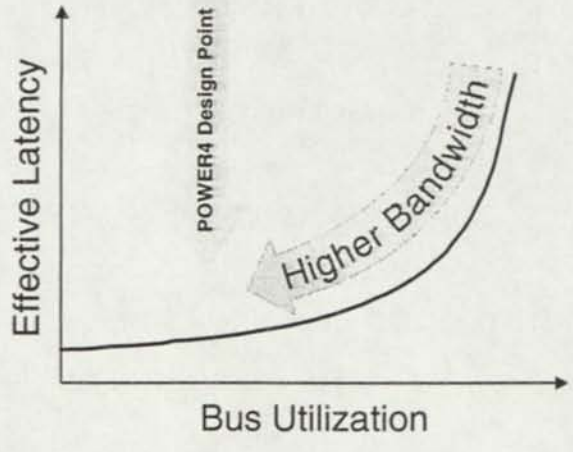

IBM POWER4

Scalability Drivers

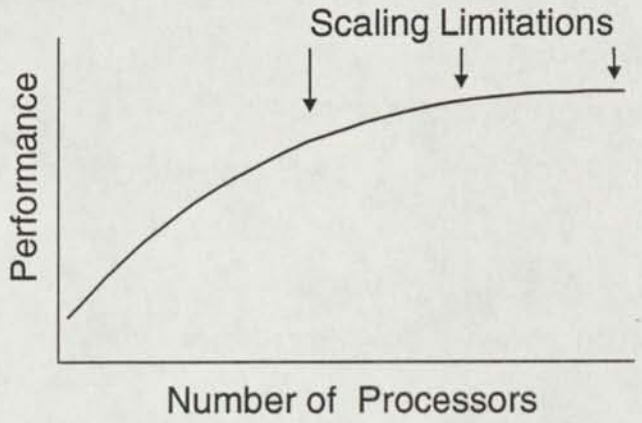

IBM

POWER4

Scalability Drivers

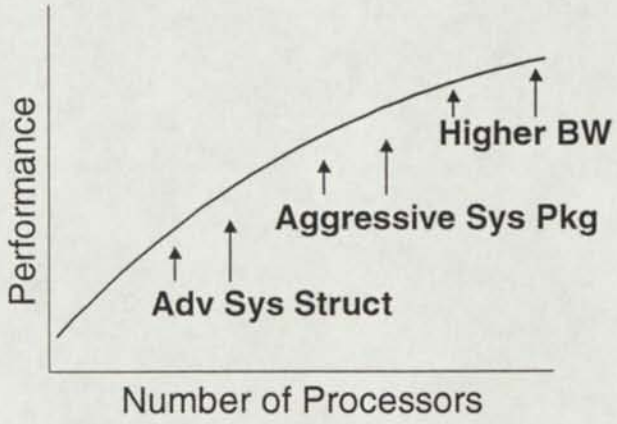

SMP performance strongly benefits
from synergistic technology deployment

IBM

POWER4

POWER4 Guiding Principles

- **SMP optimization and server flexibility**
- **Full system design approach**
- **Very high frequency design**
- **Commercial and technical workloads**
  - **Supports AS/400 and RS/6000**
- **Binary compatibility**
  - **64-bit PowerPC Architecture**

IBM. POWER4

Server Building Block
Low End Server

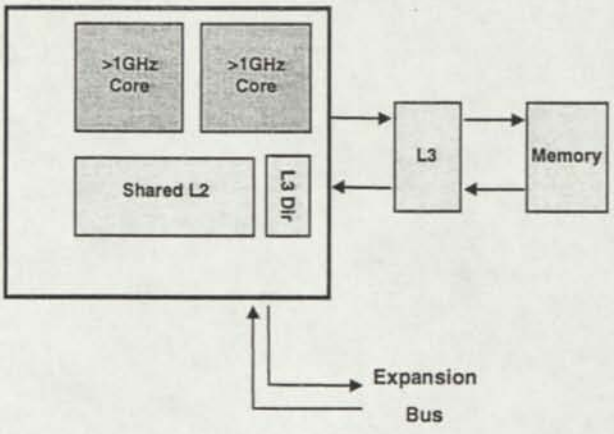

IBM. POWER4

Server Building Block

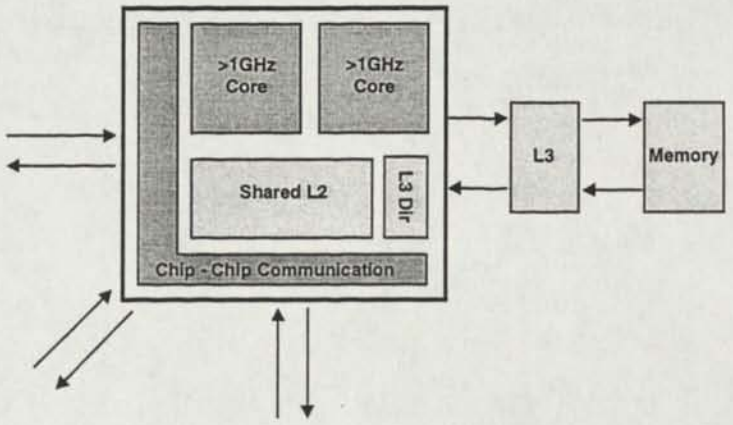

IBM

POWER4

Server Multi-chip Module

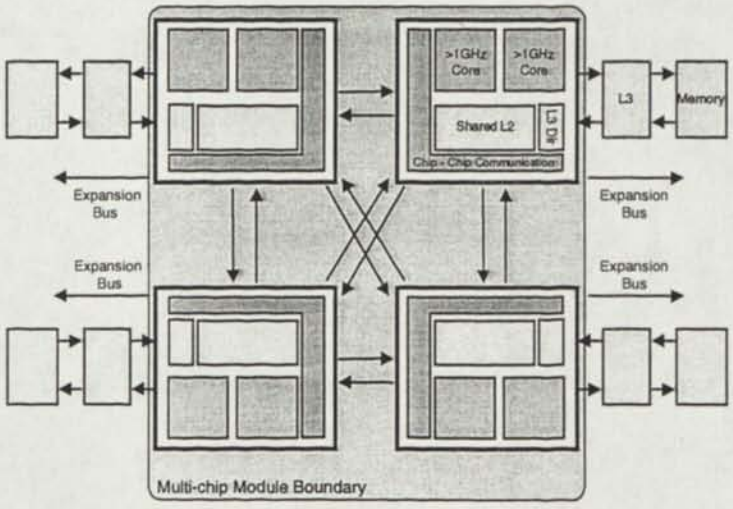

IBM

POWER4

POWER4 Memory Subsystem

- Optimized for coherency protocols
- Multi-chip module with
  - 4 chips / 8 processors
  - >100 GB/s L2 bandwidth per chip
  - >40 GB/s interface to L3 / DRAM
  - >35 GB/s chip interconnect
  - Peak bandwidths sustainable
- Expansion busses for scalability
  - SMP, NUMA, SP Clusters

IBM POWER4

Synchronous Wave Pipeline Interface

- Allows busses to operate at >500 MHz
  - 2.5-5X increase over current bus speeds
  - De-couples latency from bandwidth
  - Elastic I/O handles skew across very wide busses
- No added dead cycles required
  - Multiple requests concurrently transmitted

IBM POWER4

Chip Optimized For Frequency

- **>1 GHz chip frequency**
- **Novel machine organization**
  - **Memory latency mitigation**
  - **Hardware layering**
  - **Instruction *group tracking***
- **Large number of transistors**
  - **Data Flow design**
  - **Large number of I/Os**

IBM POWER4

POWER4 Chip Block Diagram

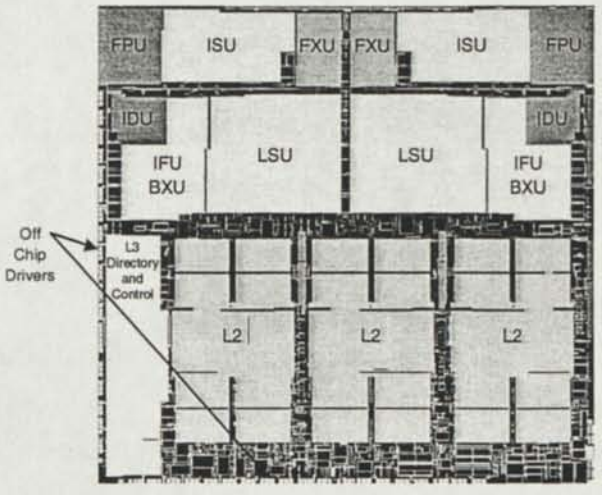

IBM POWER4

POWER4 Chip I/O Footprint

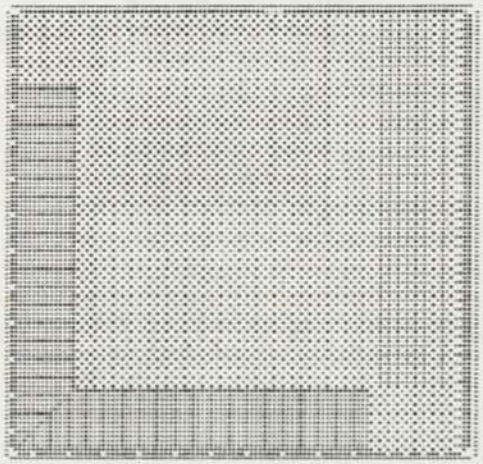

- > 2200 Signal I/Os
- > 5500 Total I/Os
- > 500 MHz Elastic I/O Busses
- > 1 Tb/s Chip Bandwidth

IBM. POWER4

POWER4 Leverages IBM Technologies

- Process
  - IBM CMOS 8S2, 0.18um
  - Copper and SOI with 7 layers of metal
  - 170 million transistors
- Package
  - Uses large number of I/Os at chip and MCM level
    - >2,200 I/O with >5,500 Pins
  - Multi Chip Module (MCM) for dense integration
- High bandwidth with fast busses
  - Elastic I/O provides >500 MHz chip-chip busses

IBM. POWER4

POWER4: Total System Design

- **Efficiently supports numerous Gigahertz processors**
  - **Integrated, distributed switch**
  - **Integrated NUMA support**
- **Performance delivered through bandwidth, frequency and technology**
- **Beginning of new era of server optimized systems**

IBM POWER4