

# Next-Generation System Design Challenges and Opportunities

By Derek Tsai

Sun Microsystems

- Contents
  - Sun's Mid-range to High-end servers and market
  - Keeping the scores – metrics of success
  - Map out the value chain
  - Identify the Bottlenecks
  - Key challenges
  - Key opportunities
  - Conclusions

# Sun's Mid-range to High-end servers: current line-up



Sun fire 3800  
Up to 8 CPUs



Sun fire 4800  
Up to 12 CPUs



Sun fire 6800  
Up to 24 CPUs



Sun fire 15000  
Up to 106 CPUs

#1 in U.S. Server Revenue (21%)\*  
#1 in Worldwide Unix Server (35%)\*

\* IDC 3/01

# Server Markets

– Fueling worldwide productivity growth

## Industries

- Manufacturing
- Telco
- Service Provider
- Retail
- Government
- Finance
- Education
- Healthcare
- Entertainment/Media

## Markets Segments

- CRM
- HPC
- Technical Development
- Collaborative
- Data Warehouse/Bus. Intel.
- Web Services
- Portal Computing
- E–Marketplaces
- Mobile
- OLTP
- ERP
- Supply Chain Management

# Server System Design Challenges - From 30,000 ft

- **Costs**
  - Competitive forces (dot com burst, Wintel servers)
  - Perception of Benefits (Wintel servers, bad economy)
- **Thermal/Power management:**
  - From the box out (2x to 3x of power density/generation)
  - From the data center in (Max'ed out on cooling, power shortage)
- **Scaling Up of performance**
  - Stressing interconnect bandwidth
  - More modules → more interconnects
- **Boosting RAS (Reliability, Availability, Serviceability)**
  - Complexity ↑, Heat ↑, Bandwidth ↑, Interconnects ↑

# Keeping the scores – metrics of success

- Why keep scores/metrics?
  - Too easy to over-design and add unnecessary costs
  - De-mystify the signal-integrity black magic
  - Capture Signal-Integrity Engineers' value-add (not a prohibitor)
  - Know the trend to reduce risks
  - Sun's Six-sigma initiative (Sun Sigma)

## Keeping the scores – metrics of success (Cont'ed)

- Possible metrics:
  - I/O Bandwidth
  - I/O power consumption
  - I/O noise margin
  - Power noise
  - Pin/route density
  - Power density

## Keeping the scores – metrics of success (Cont'ed)

- Ultimate benefit for systems
  - I/O bandwidth/throughput
  - Latency (dominated by placement, architecture, and speed of light)
- Costs
  - \$ (component and implementation)
  - power/thermal
  - Noise (electrical and Electromagnetic)
  - Packaging (pin/route density, material, & etc.)
  - Reliability

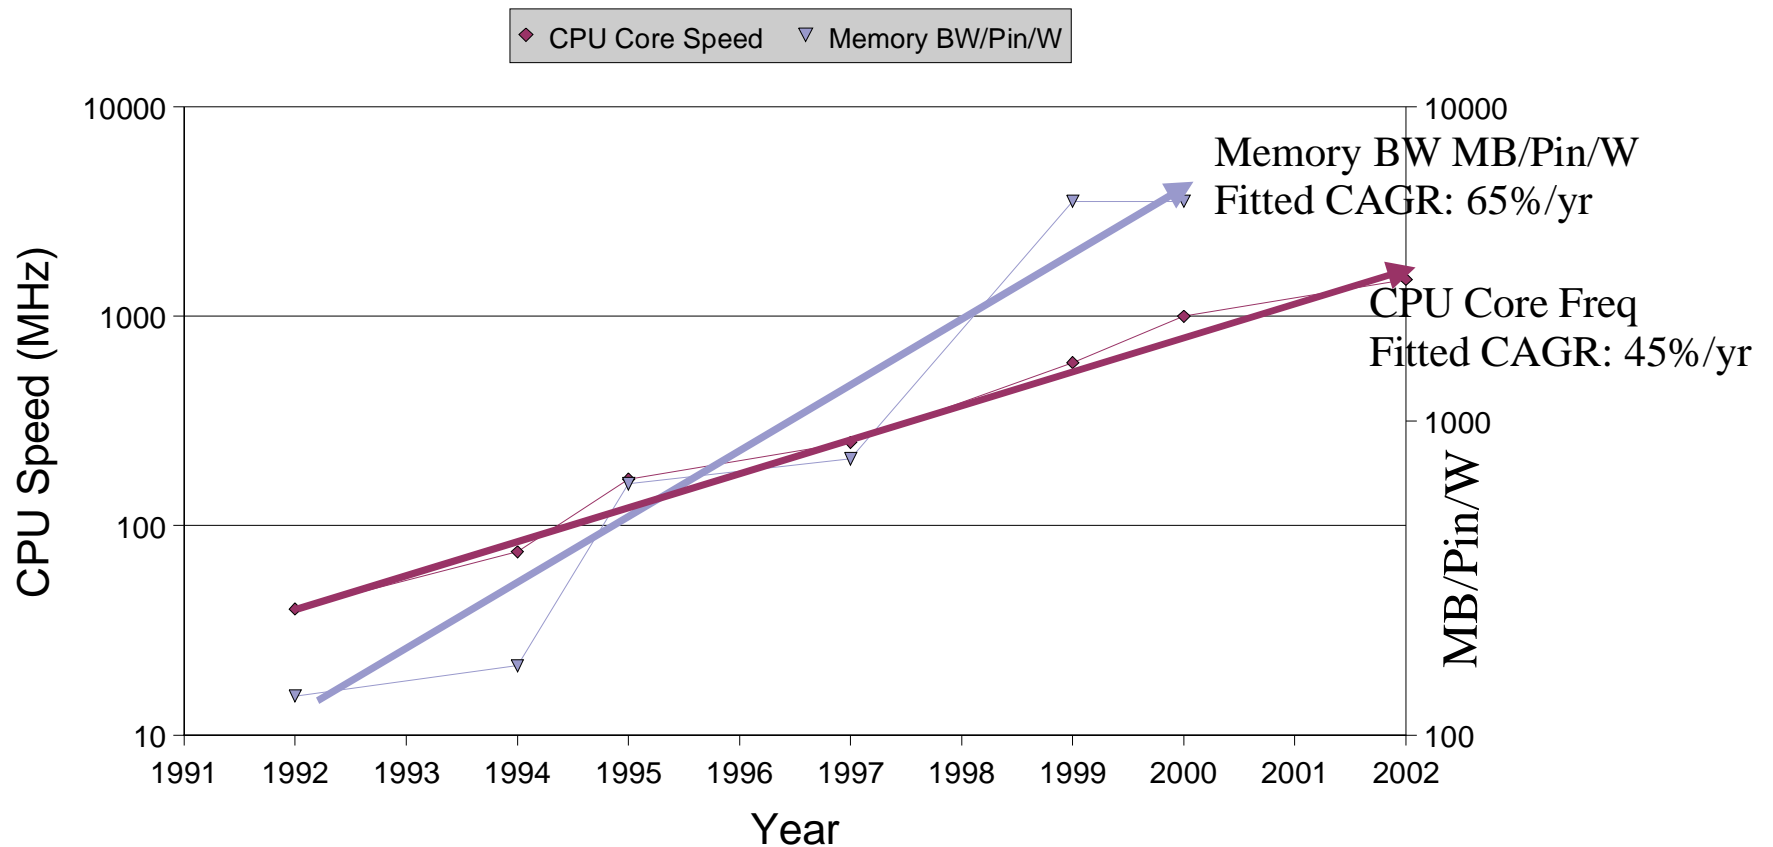


## Keeping the scores – metrics of success (Cont'ed)

- Proposed metrics for SI/Packaging Engineers:
  - I/O bandwidth / pin / **Watt consumed** / \$ AND
    - Meeting spec'ed reliability
    - Meeting EMI regulation
    - Having sufficient noise margin for all component and manufacturing process corners
  - $\Delta\%$  Operating Freq Margin /  $\Delta\$$ 
    - **AND** same conditions as above

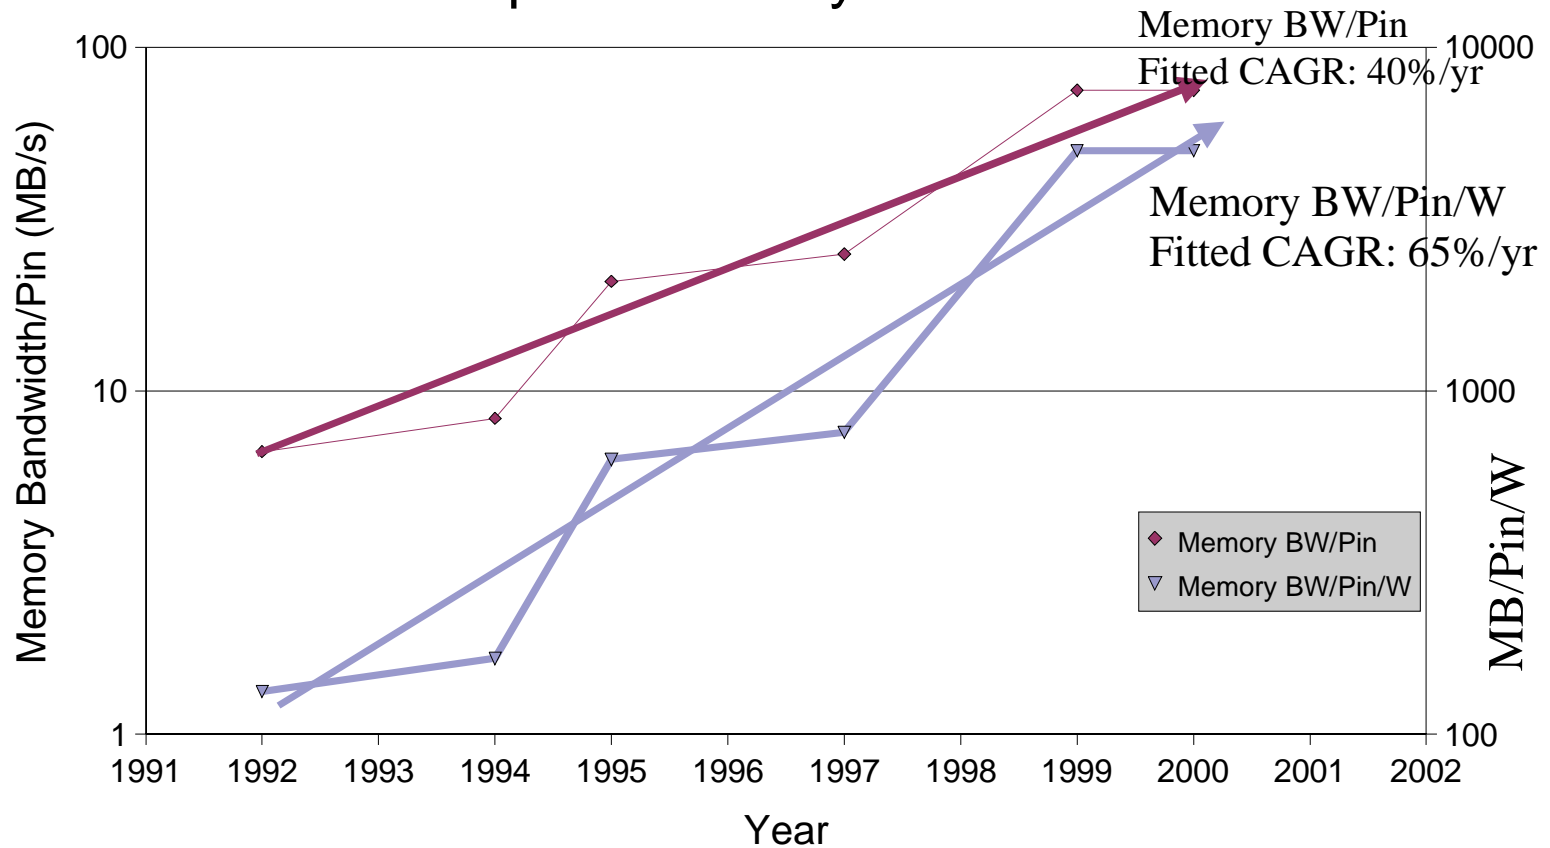
# Keeping the scores – metrics of success (Cont'ed)

Sun's Sparc CPU Speed & I/O Bandwidth



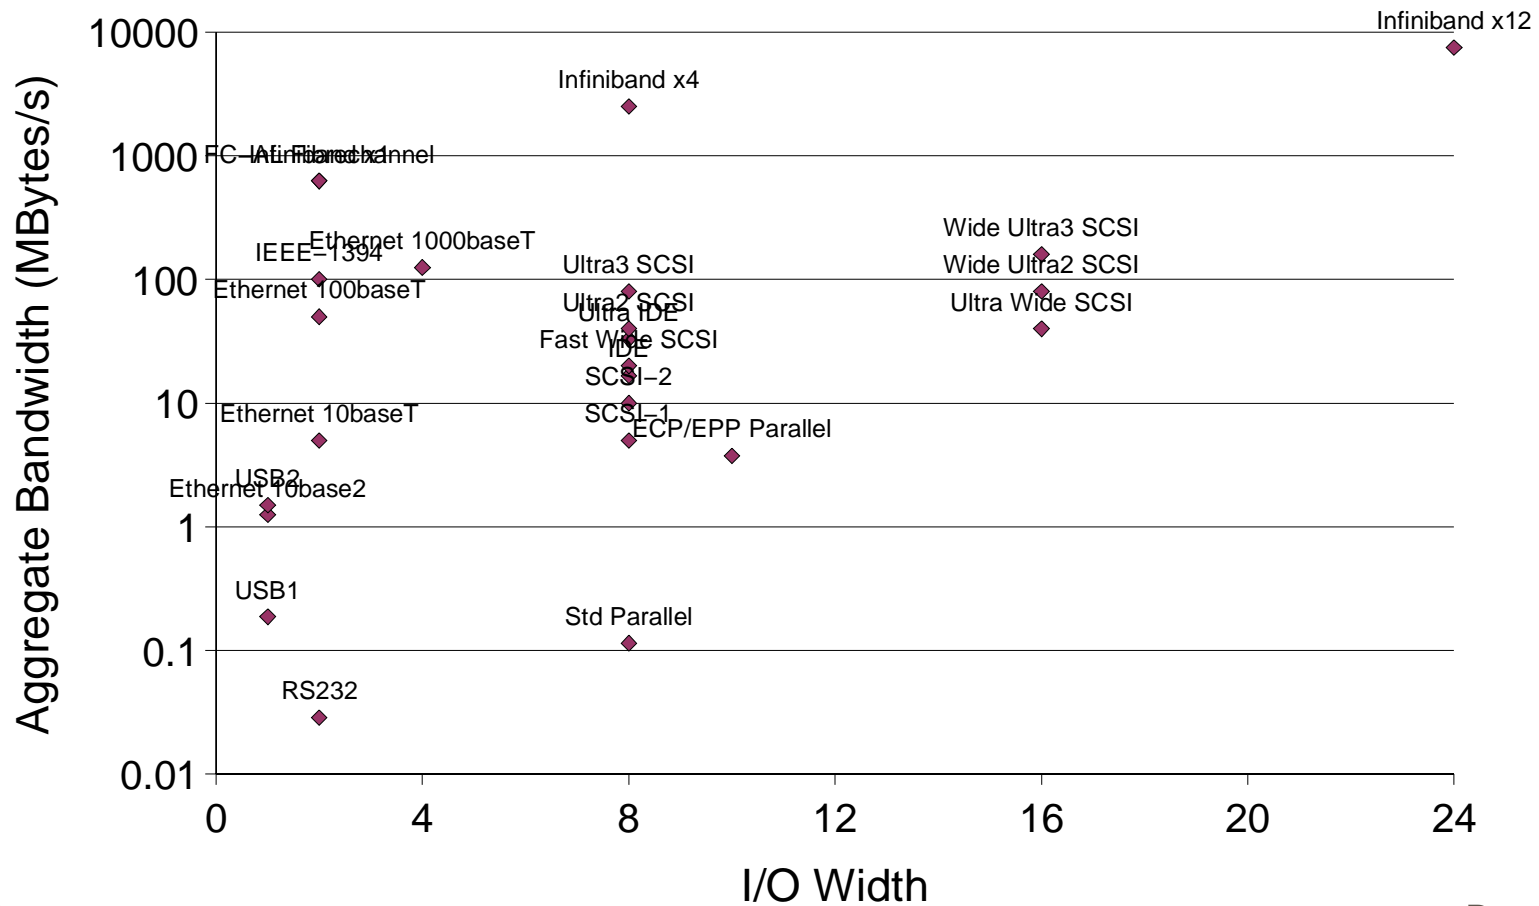
# Keeping the scores – metrics of success (Cont'ed)

## Sun's Sparc Memory I/O Bandwidth



# Keeping the scores – metrics of success (Cont'ed)

## I/O Standards Survey

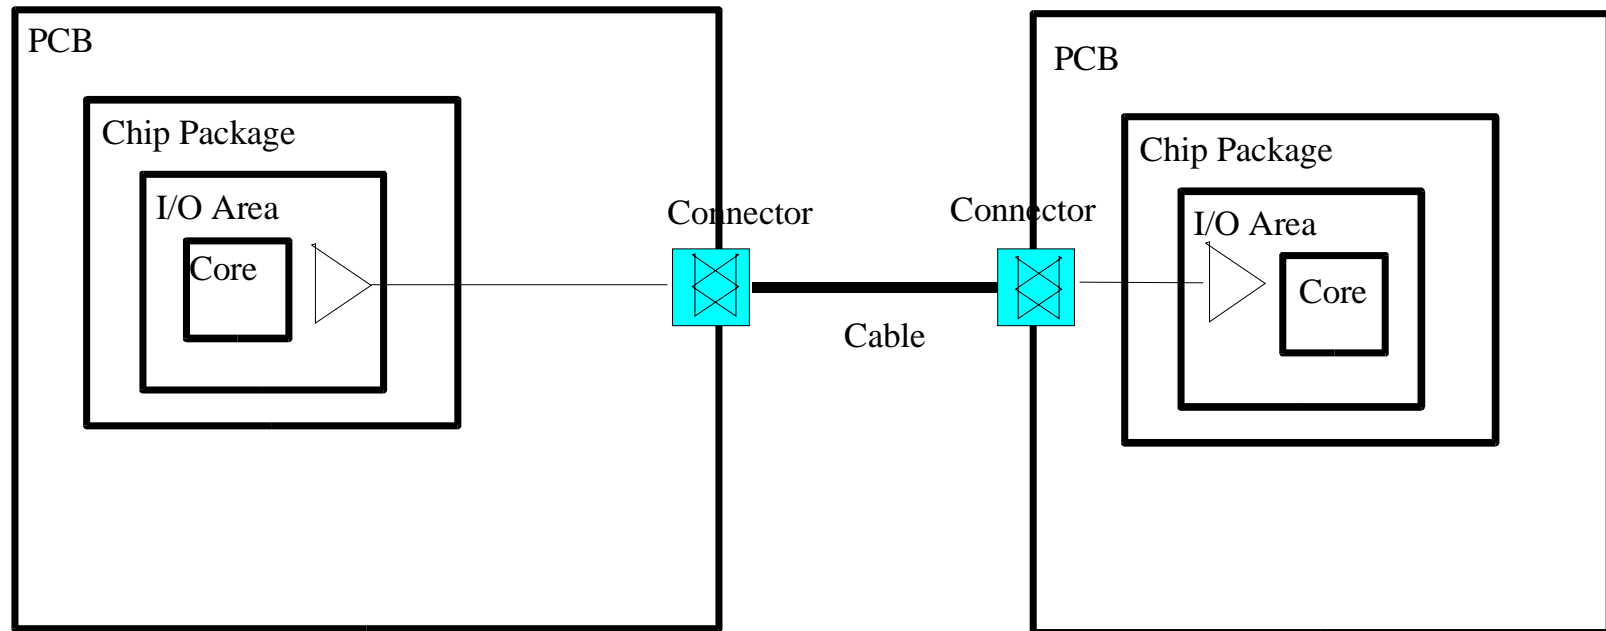


# Mapping the value chain of system design

- Why map the value chain?
  - Determine the bottleneck and invest to remove
  - Bottlenecks move around for each generation
  - Reduce the risks
  - Reduce the costs of over-design
  - Because it's a SYSTEM

# Mapping the value chain of system design

## System Level: Signal Path



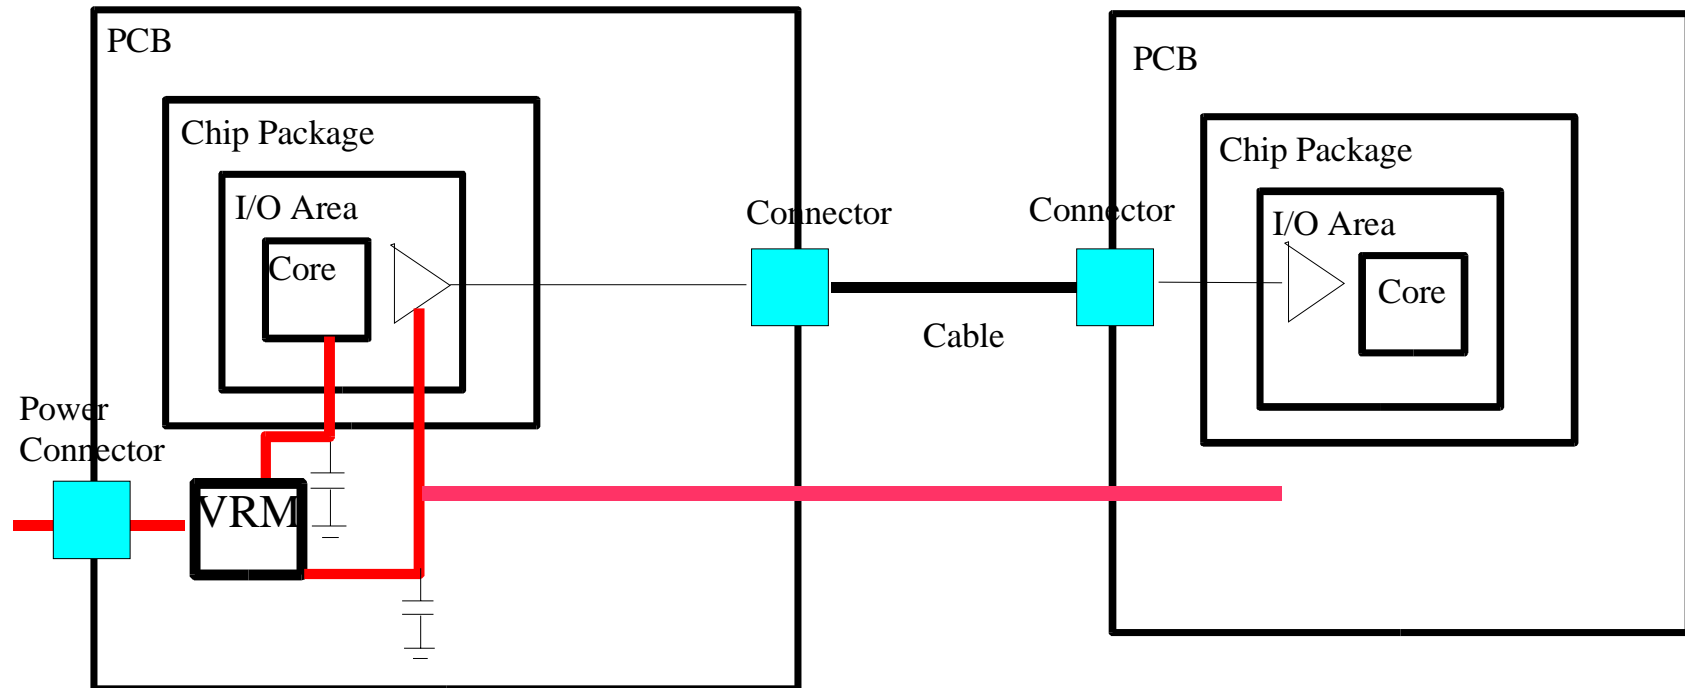
I/O Cell, ESD , Chip package, PCB trace, vias, connector, transceivers, cables

# Identifying the bottlenecks

- Signal path bottlenecks
  - I/O Cell Design: from Global Sync to Source Sync
  - Simultaneous switching I/O noise
  - ESD load
  - Package trace
  - Package/PCB interface
  - PCB trace loss and routability
  - PCB vias
  - Signal reference
  - Transceivers
  - Cables (copper and optic)

# Mapping the value chain of system design (Cont'ed)

## System Level: Power Path



Power connector, VRM, Capacitors, PCB planes, on-package caps, on-die caps



# Identifying the bottlenecks

(Cont'ed)

- Power path bottlenecks:
  - Power connector – inductance & resistance
  - VRM – power transient, density
  - PCB Power – spreading inductance
  - Package Power
    - Core Power – power transient, package resonance
    - I/O Power – return reference, SSN
  - Power across domains – cabling

# Key Challenges

- Interconnect density
  - Mechanical:
    - Qualification of connectors, pcb and packages
    - Finding the failure modes (too many interfaces)
    - Assembly and handling
    - Routing and pin escapes
    - Lead-free
  - Electrical:
    - Maintaining the signal quality: lossiness, via loss, crosstalk, return paths, SSN, Bit Error Rate, and etc.
    - EMI
    - Diagnostics and testability
    - Observability: design for debug and measurement
  - Economics: total costs of ownership/design

# Key Challenges (Cont'ed)

- Power and Cooling
  - Delivering power:
    - Power architecture (staging/budgeting) critical
    - Optimal use of decoupling capacitors: VRM, board-level, package-level, and chip-level.
    - Availability and accuracy of capacitor models
  - Thermal:
    - Efficient use of power – justify the demands
    - E star mode

# Key Challenges (Cont'ed)

- Design Complexity
  - Model availability and accuracy:
    - SPICE vs. IBIS vs. others: methodology wanted
    - Frequency–dependent interconnect models
    - Resource drain to provide accurate models
  - Simulation:
    - HSPICE vs. SpecctraQuest vs. XTK
    - What to look for: timing, error rate, eye diagram
    - What to include in simulation: – too many variables, information overload – determine what are essential.
  - Tedious – must deploy computer tools/scripts to excel

# Key Challenges (Cont'ed)

- SI Engineering Resources
  - Lack of hands-on debug experience: video game culture creates lots of naïve engineers – lack "instincts"
  - Specialized to be competent, but not too specialized to be practical – insufficient "system" view
  - No glory: hard to extract our value-add – need good success metrics
  - Case studies needed in teaching – SI no less complex than business cases

# Key Opportunities

- Move away from Black Magic to Standard SI design and test methodology: put aside the ego and show me the data
- Industry-wide Simulation methodology: Component suppliers to provide accurate models, working with tool providers
- Industry tool certification body (JEDEC) supplying benchmark test cases like TPC for database
- Specialized SI curriculum with lots of real-life case studies provided by industry

# Conclusions

- Interconnect/Packaging is the bottleneck now and System is as good as the bottleneck: must solve it as a system to avoid large \$ and risk costs
- Some of Sun's challenges are shared by the industry:
  - Better metrics of success – common language/goal
  - Joint solution can enlarge the pie for all and start with providing/sharing accurate models, and
  - Open interface to SI tools allows users to tackle next design challenges