IBM Research

Future Design Systems for Large Chips

John Darringer
IBM T J Watson Research Center

Outline

- **Trends**
  - Business
  - Technology
  - System
- **Impact on Design Systems**
- **Summary**

Business Trends

- **Information technology business is growing again**
- **Companies re-engineering their business**
  - Pervasive networking
  - Ubiquitous sensors
  - Affordable computation
  - Open standards
- **Innovative solutions provide completive advantage**

Growing Demands on System Designer

- **Exploit**
  - Increasing parallelism
  - Low latency communication
  - Accelerators for critical tasks
- **Offer**
  - Scalable systems for growth
  - Reliable services and storage
  - Affordable cost of ownership

Innovative Architecture Example 1 – Cell Processor

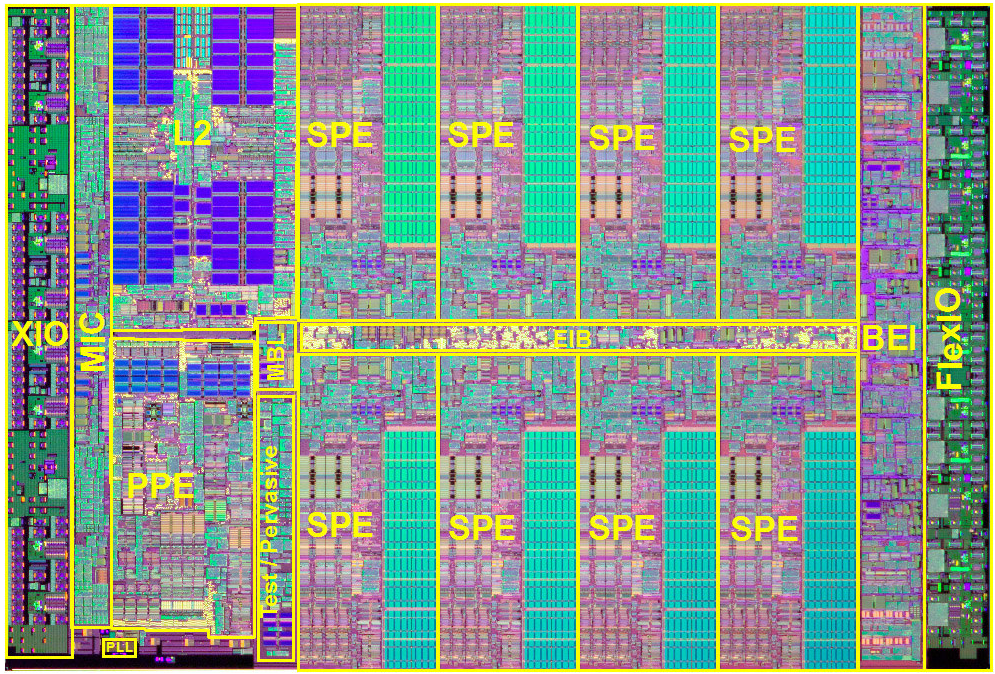

- **IBM, Sony and Toshiba partnership**
- **64-bit Power Architecture™ (PPE)**
- **8 Synergistic Processor Elements (SPE)**

Innovative Architecture Example 2 – Blue Gene/L

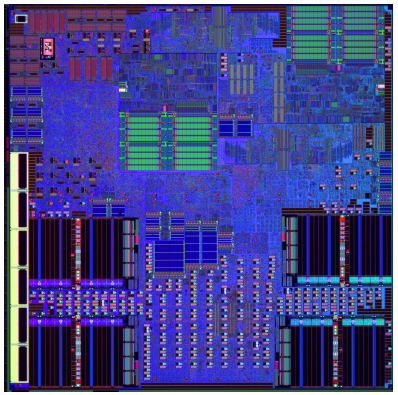

700 Mhz SoC

Chip
2 processors

2.8/5.6 GF/s
4 MB

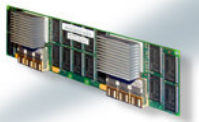

Compute Card
2 chips

5.6/11.2 GF/s
1.0 GB

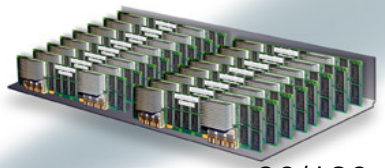

Node Card
32 chips

90/180 GF/s
16 GB

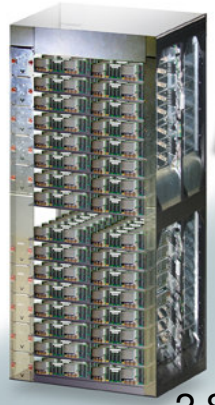

Rack
32 Node Cards

2.8/5.6 TF/s
512 GB

System
64 Racks, 64x32x32 chips

180/360 TF/s
32 TB

System Design Becoming More Complex

- **Complex tradeoffs – power, performance, interconnect**
- **Evolving and future workloads**
- **New software architectures**
- **Done by a few experienced experts**
- **Tools are spreadsheets and custom models**

Need to Extend DA to System Design Craft

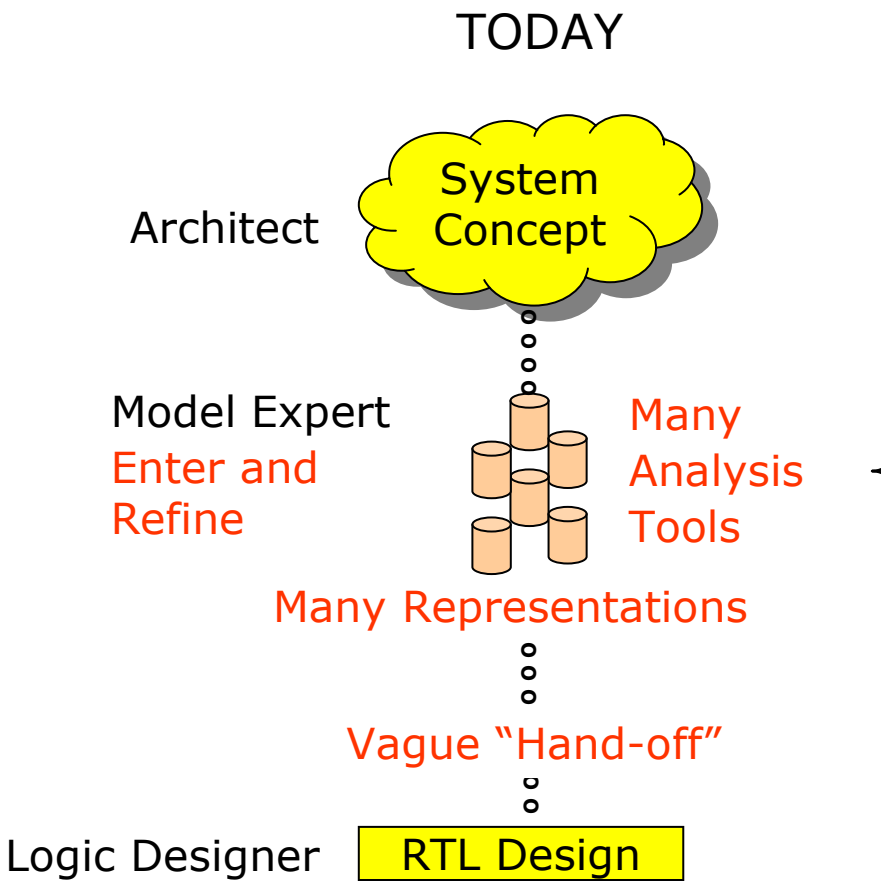

- **Spreadsheets**
- **Technology Models**
  - **Energy, power, temperature**
  - **Area, yield**
  - **Package**
- **Performance Models**
  - **Core**
  - **Cache**
  - **Accelerator**
  - **Bus**
  - **Memory hierarchy**
  - **I/O network**
- **Workload Traces**
- **Each with unique input and outputs**

Growing Technology Challenges

Performance

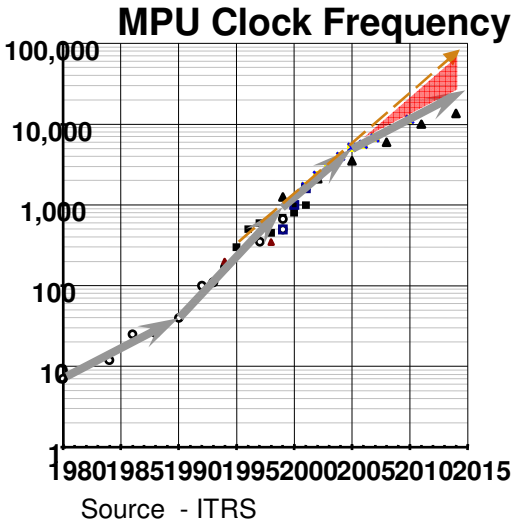

Power Management

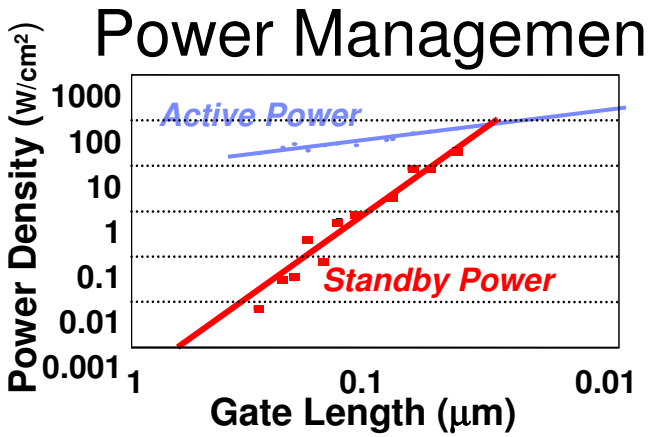

Variability

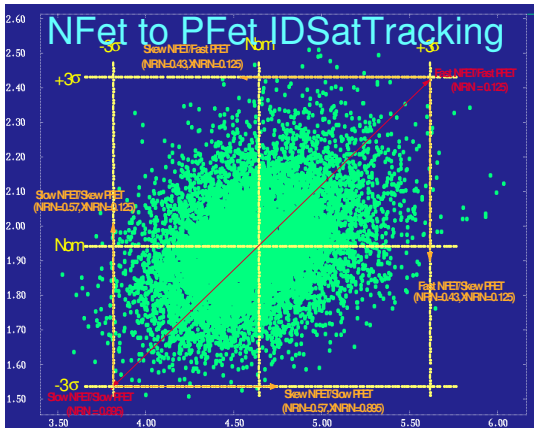

Manufacturability

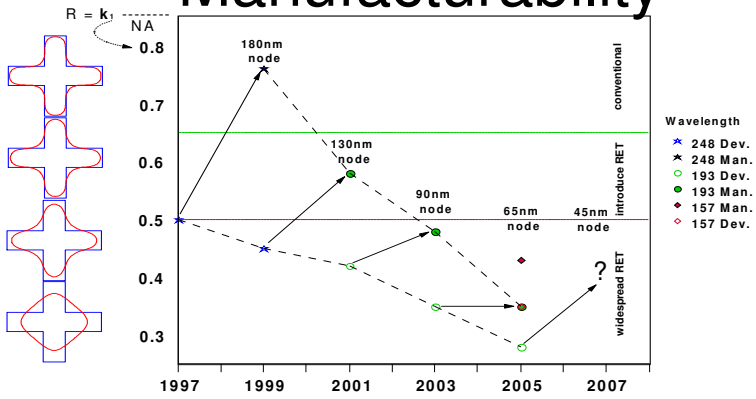

Time

Systems Performance – An Integrated Approach

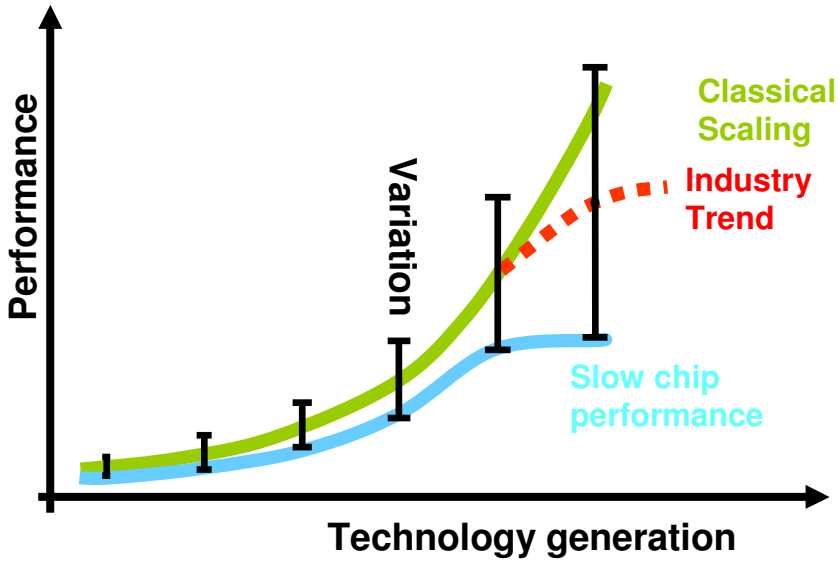

- **Frequency no longer dominant performance driver**
- **Advances will come from entire performance stack**

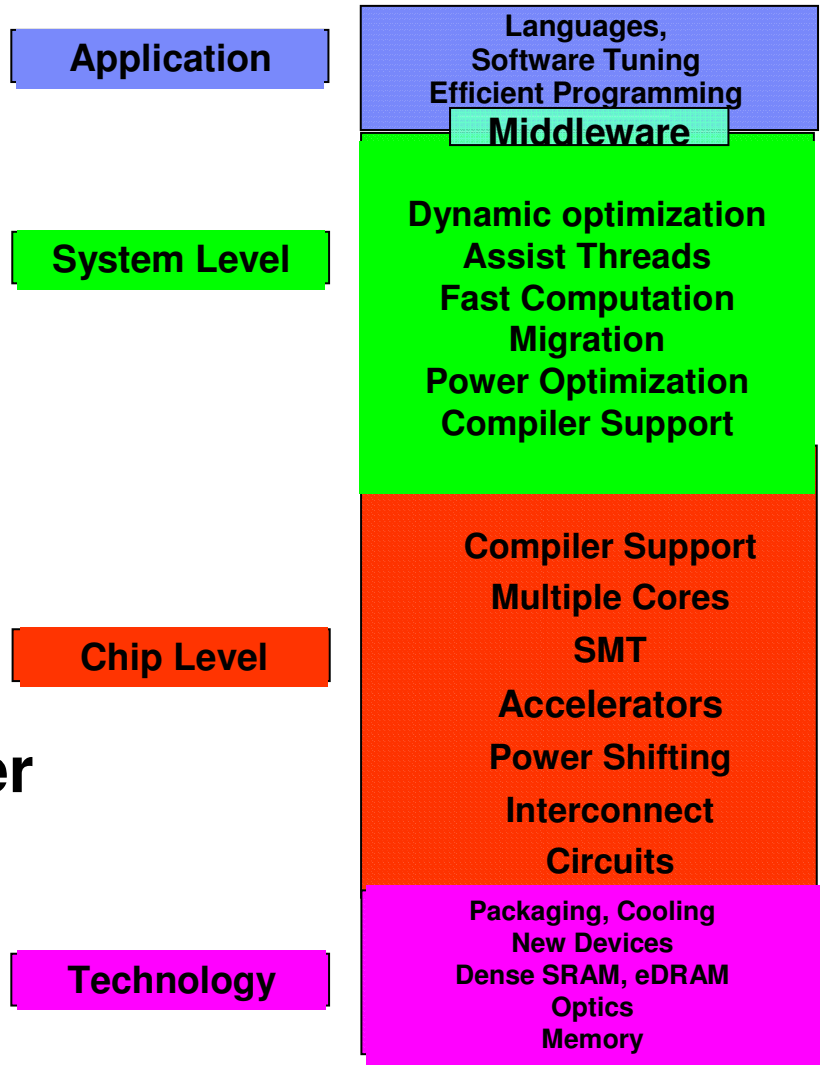

Complex Hardware Tradeoffs

- **Multi Vt, clock gating, power gating**
- **Voltage islands, voltage and frequency scaling**
- **Asynchronous design**
- **Multiple cores**
- **Accelerators**
- **Memory hierarchy**
- **I/O network**
- **New packaging concepts**

Example - Multiprocessor Accelerator for SoC Design

- **Self-contained, scalable multiprocessor core**
  - With its own general-purpose processors, interconnect, interfaces and memory
- **Connected to the SoC bus via a bridge**
  - Easy integration in the basic SoC structure
- **Inner structure and complexity is hidden from the SoC designer**
  - Significantly simplifies the SoC design

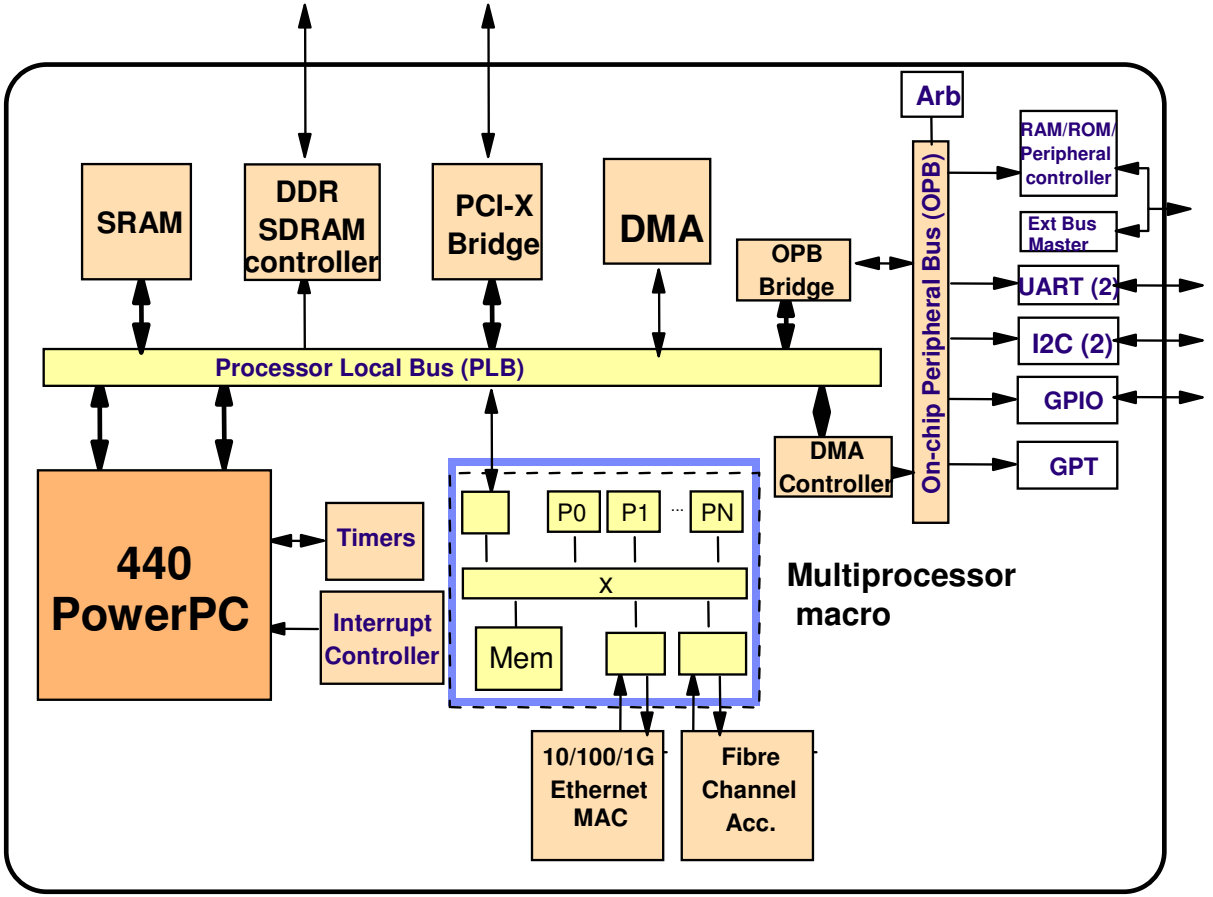

Accelerator Performance and Floorplan

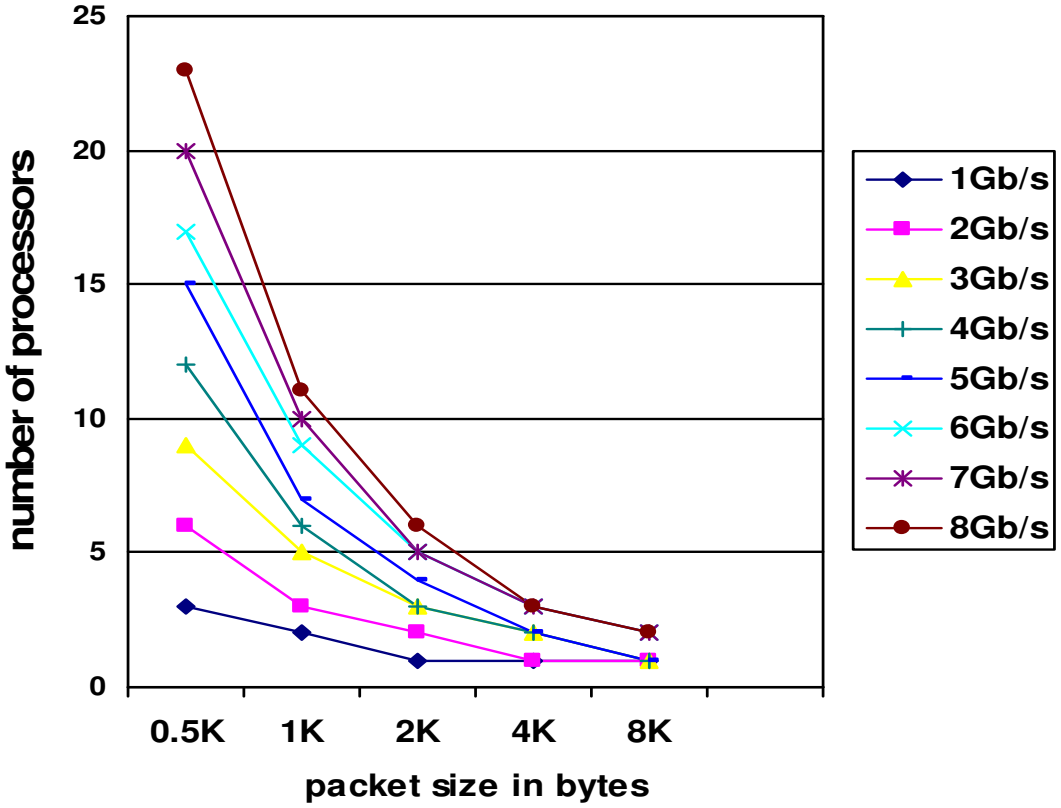

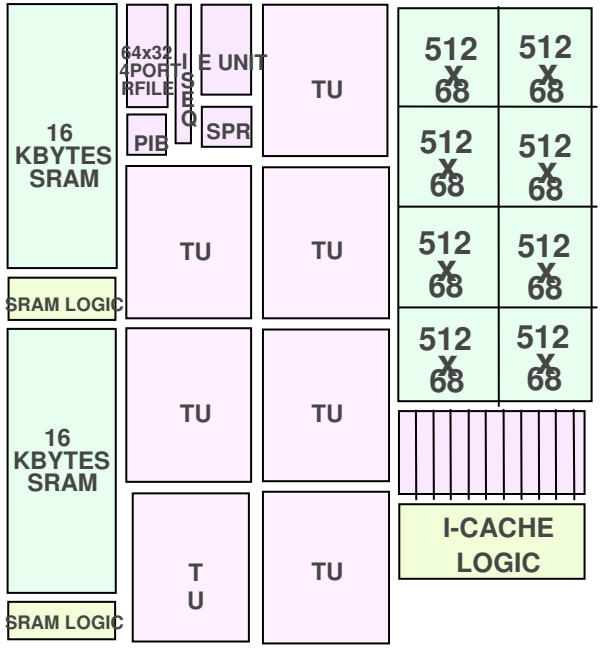

- **Parametric core offers wide range of performance**
- **Need to evaluate novel architectures in system context**

Globally Asynchronous Locally Synchronous - GALS

- **Voltage islands enable multiple frequency domains**
- **How many partitions?**
- **Where is best place to partition?**
- **What is synchronization penalty?**
- **Need flexible model to explore sensitivities**

GALS Example - Out of Order Machine

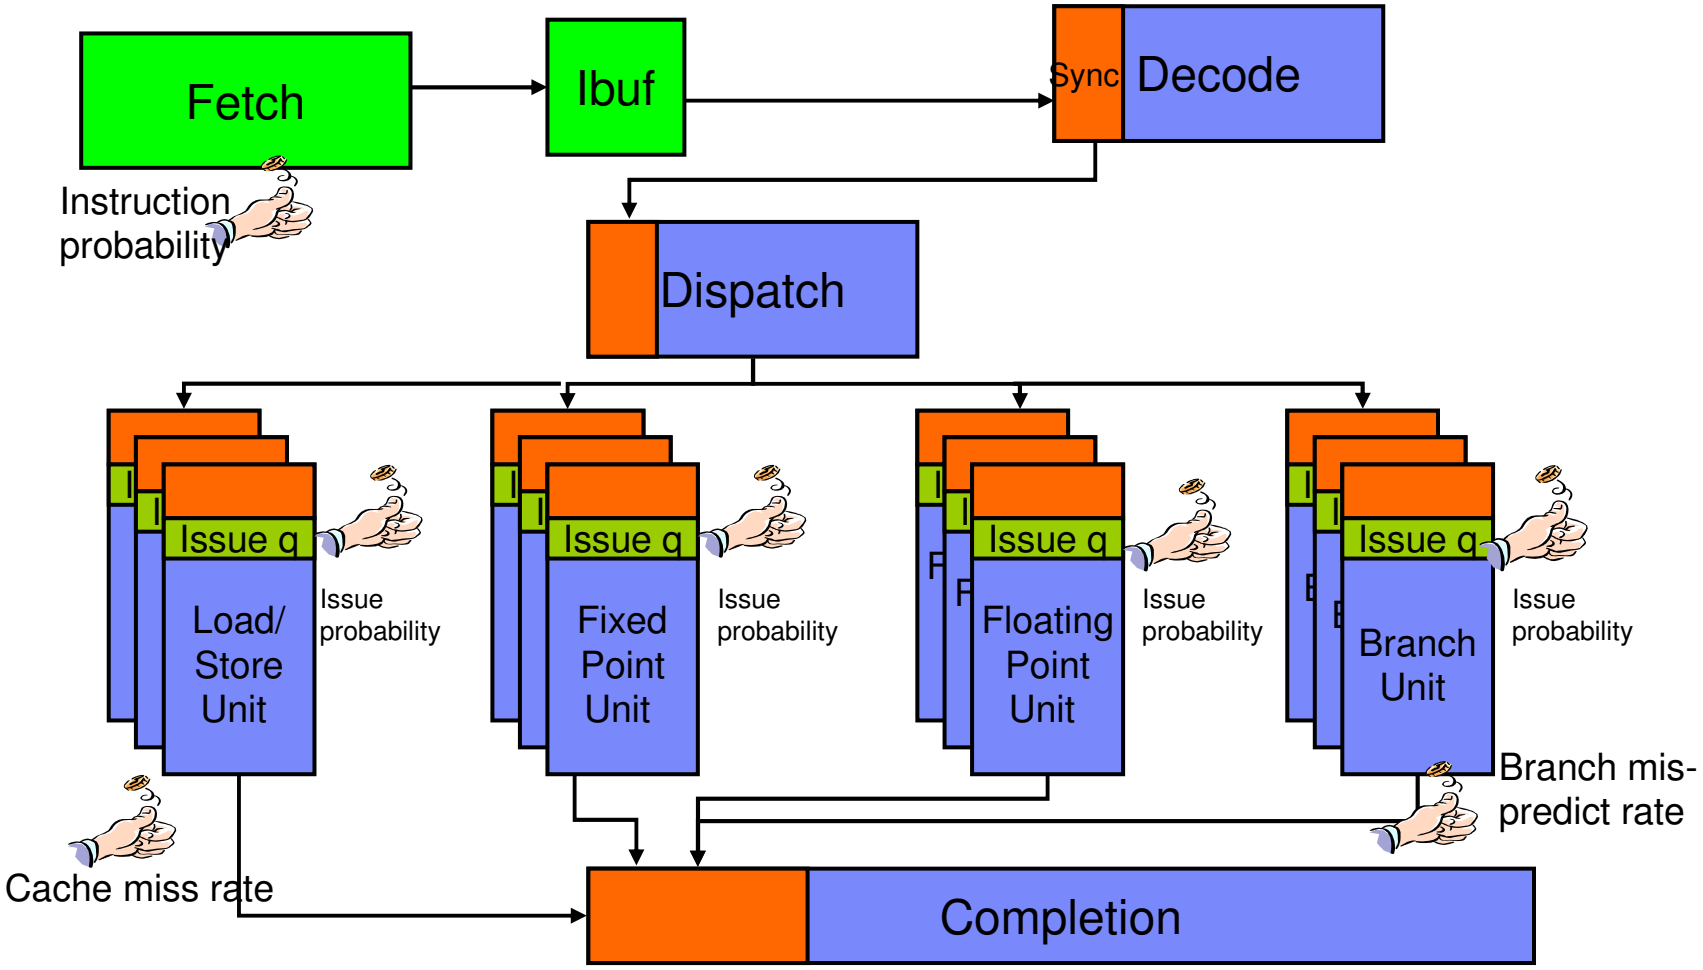

One GALS Experiment - Half the Fetch Frequency

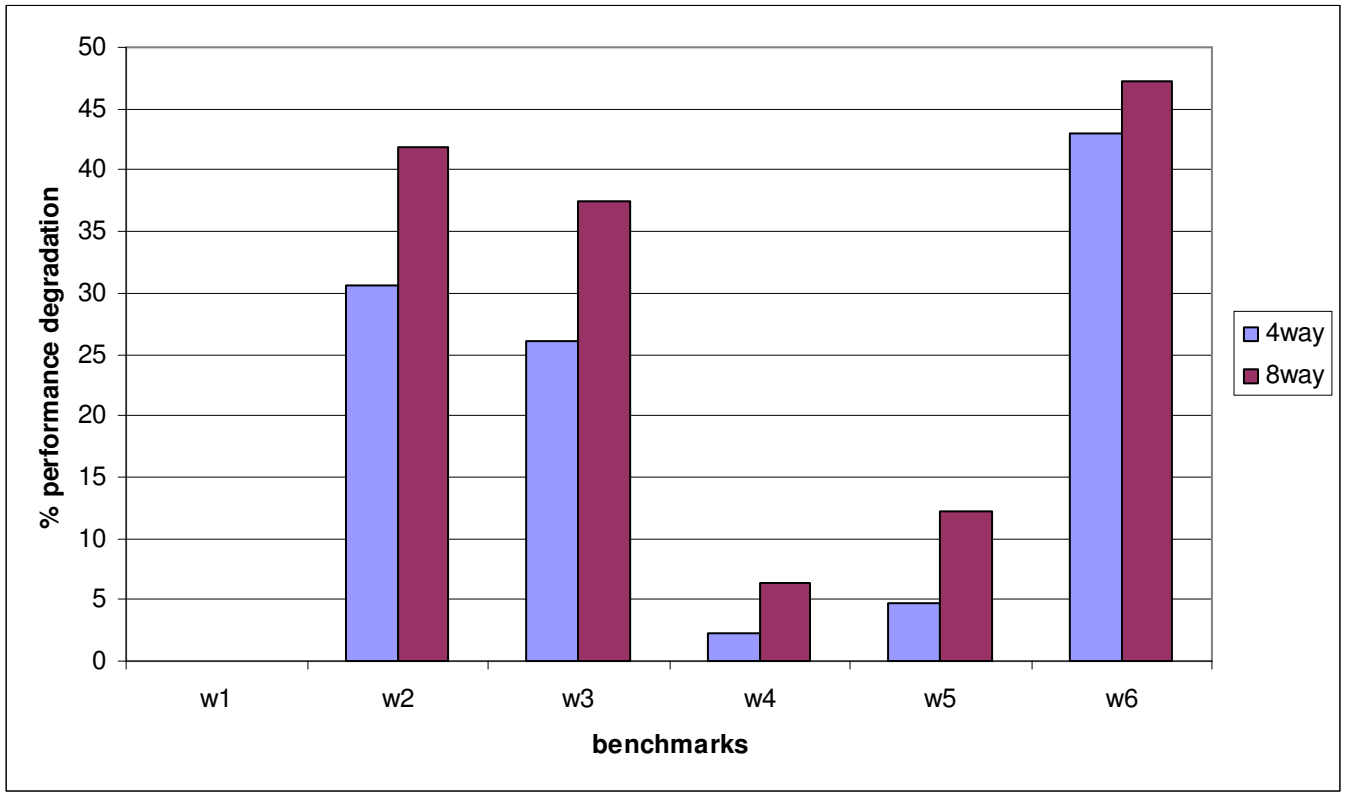

- Out-of-order machine, fetch clock 20ns, others 10ns, no sync penalty
- Compared with all units clocked at 10ns

Need Tools to Evaluate Tradeoffs Early in Design

- **Intuitive design capture**
- **Unified representation above RTL**
- **Rapid and accurate tradeoff analysis**
  - Performance, power, interconnect
  - Handle technology options
    - Voltage Islands, multi-Vt, frequency scaling, clock gating
- **Linkage to RTL flow**

SEAS - SoC Early Analysis System

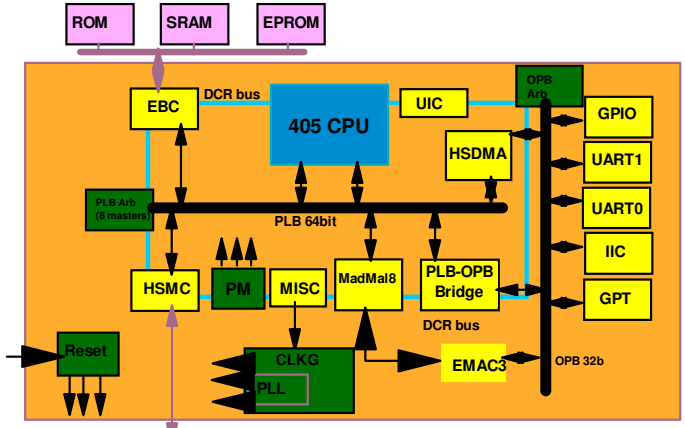

SoC Block Diagram

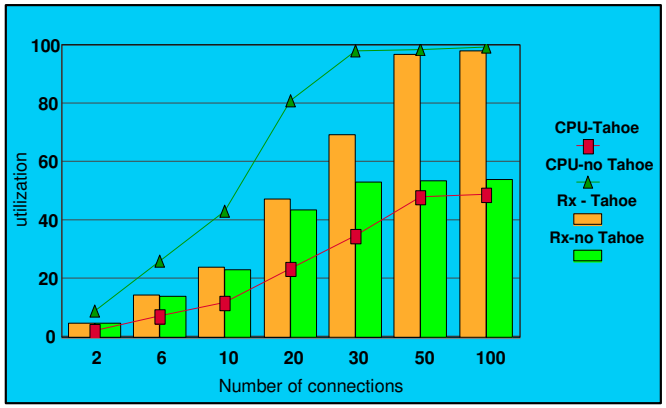

SoC Performance

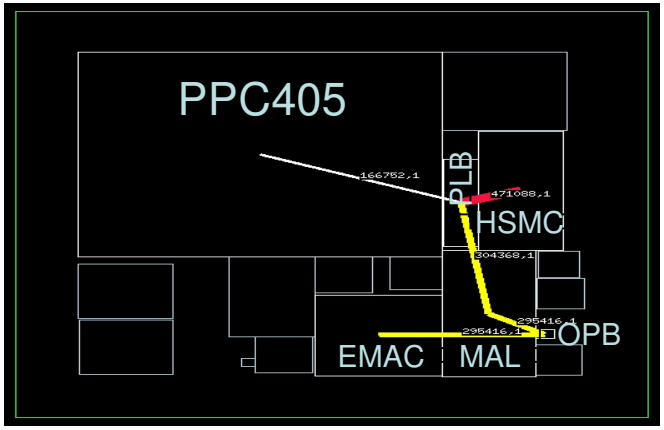

SoC Floorplan

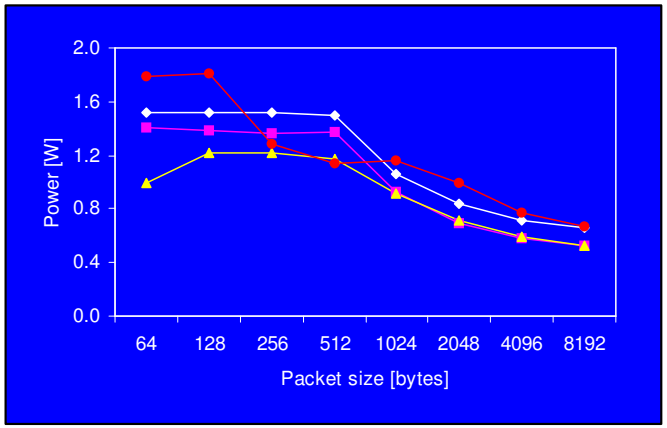

SoC Power

SEAS Environment

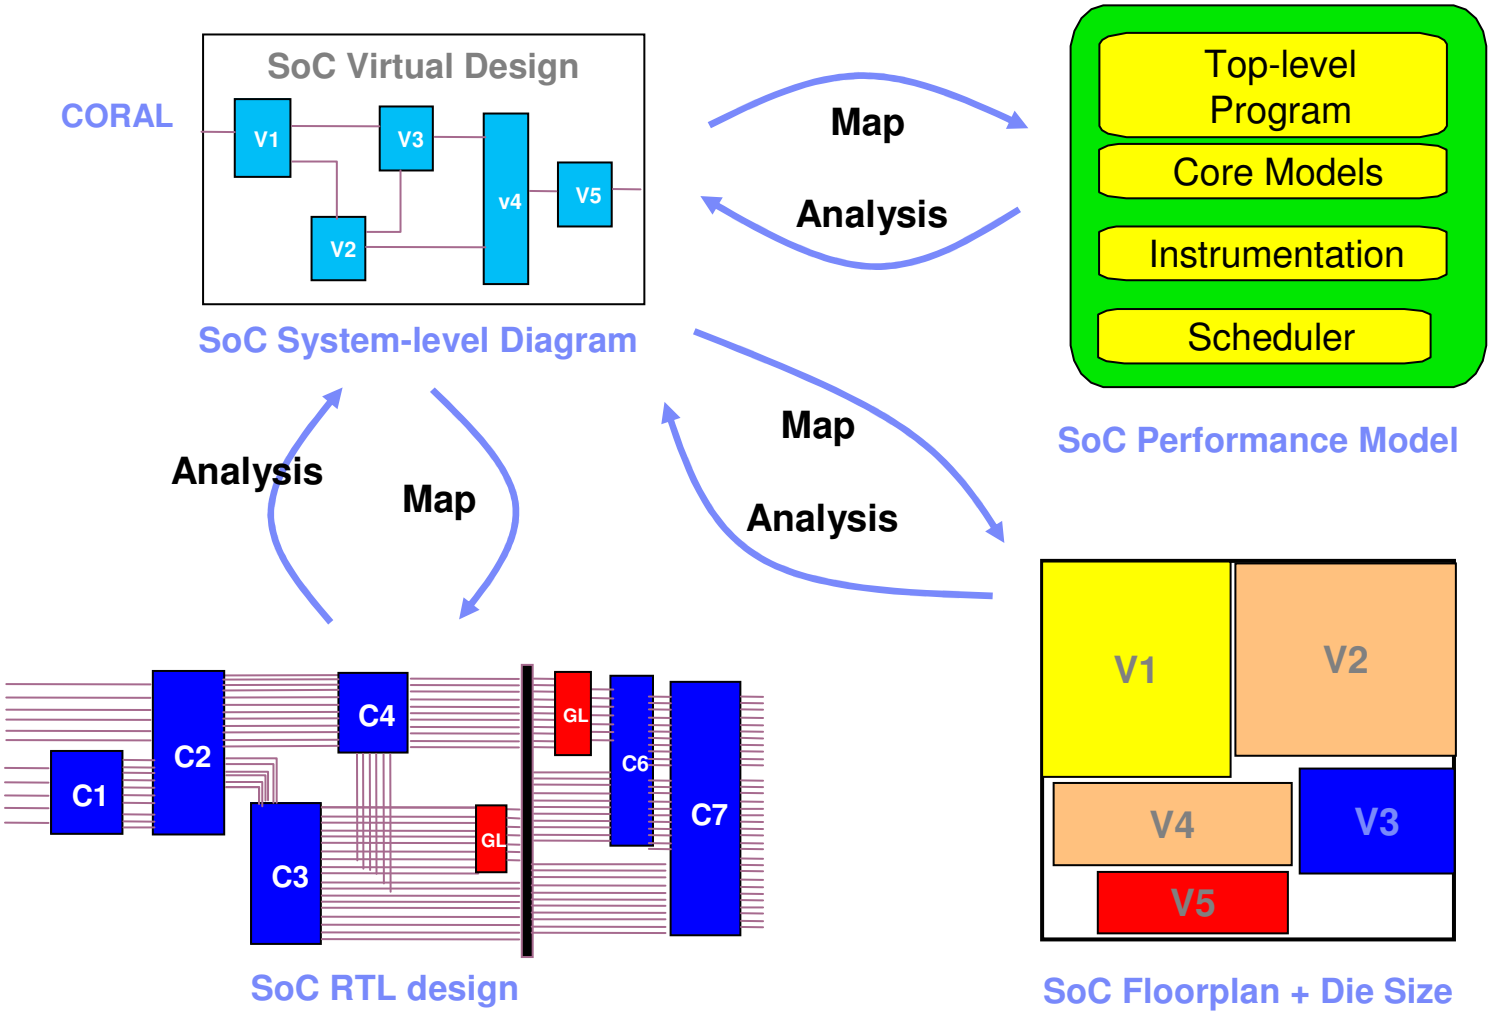

SEAS Experiment

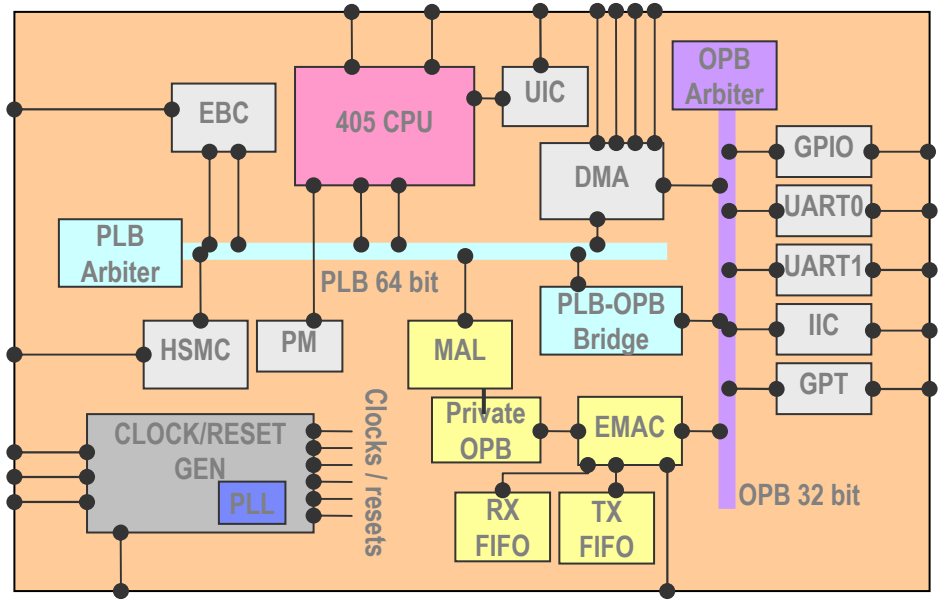

- **405 Platform**
  - Ethernet Subsystem
- **Add cores for capacity**
- **Assess Impact**
  - Performance
  - Power
  - Chip size

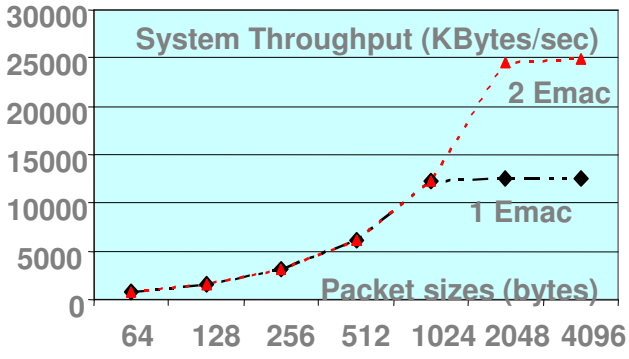

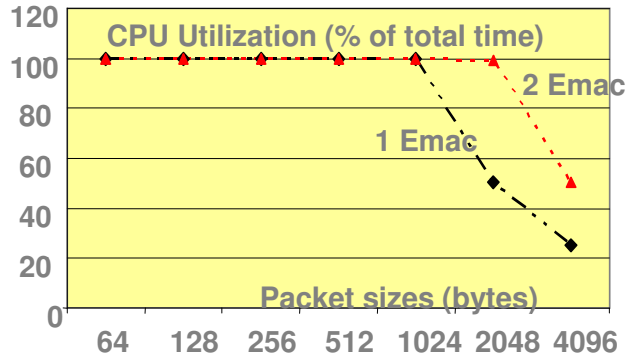

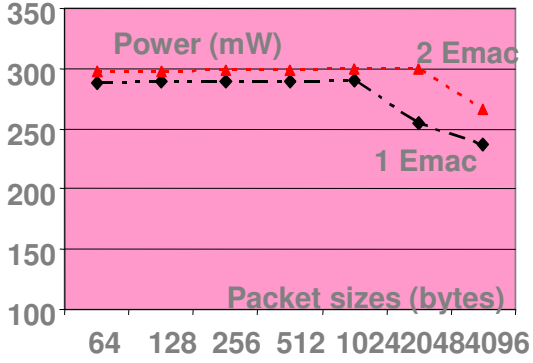

SEAS Experiment

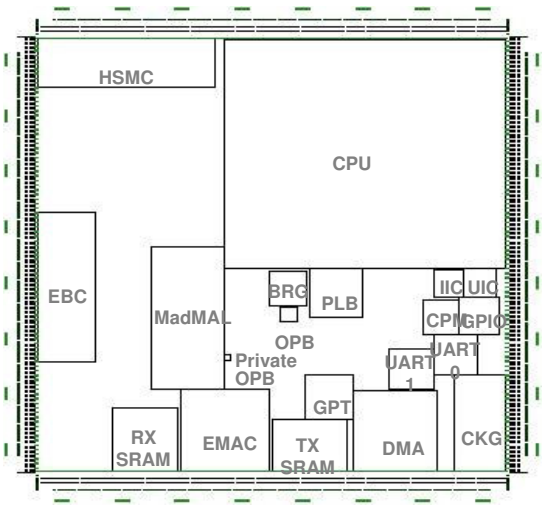

1 EMAC
6.05mmx6.05mm

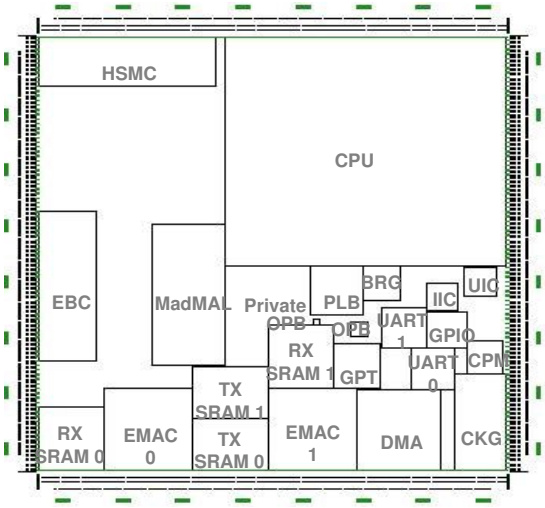

2 EMACs
6.05mmx6.05mm

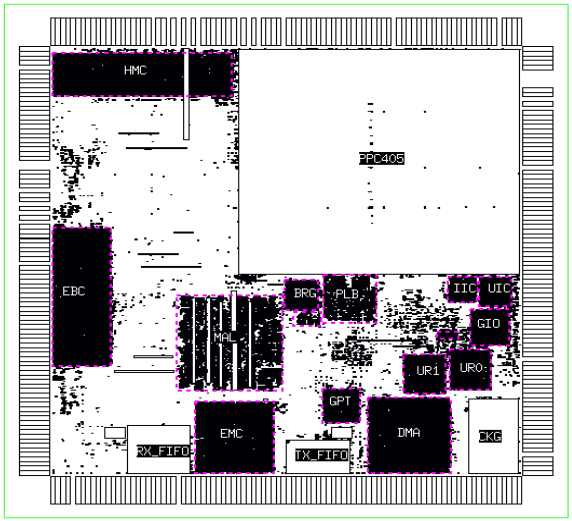

Real Design Layout (1 EMAC)
6.05mmx6.05mm

- **Results**
  - Revised design met performance requirement
  - Could fit in original die-size (pin limited)
  - Floorplanned implementation met timing requirements

System Trends

- **Today's ASIC and processor design methodologies are distinct**
- **As technology enables greater integration they will converge**

IBM ASIC Characteristics

- **Large chips**
  - 2X industry average
- **Rapid TAT**
- **First time right**
- **SoC approach**
  - Pre-verified reusable cores
  - Design at system level
  - Rapid implementation

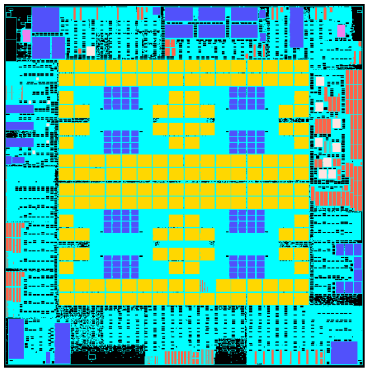

40-Terabit Router Engine

- 38 M gates
- 130 nm
- 18.3 mm x 18.3 mm
- 7 Levels of Metal
- 3389 Total I/O

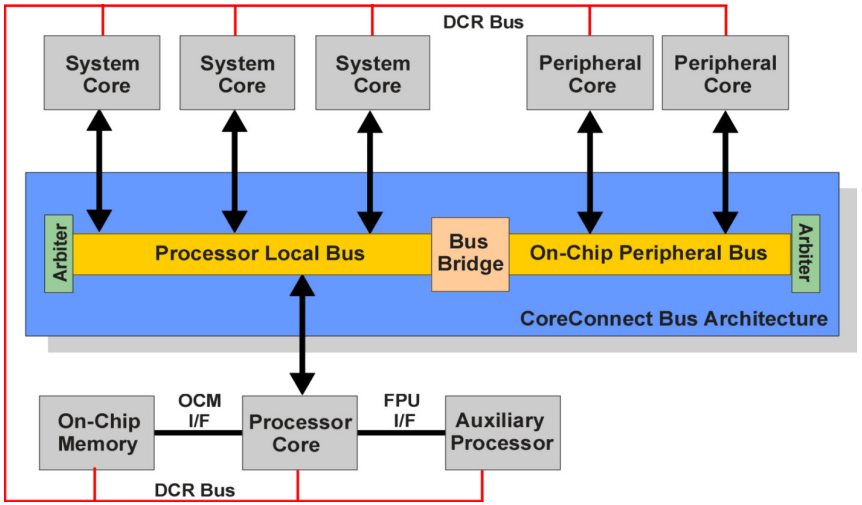

IBM Core Connect Architecture

IBM Processor Characteristics

- **Large chips**
- **Custom macros**
- **50% synthesized logic**
- **First to see challenges**
- **First time right**
- **Multi-core**
  - Increasing parallelism
  - More moving on chip

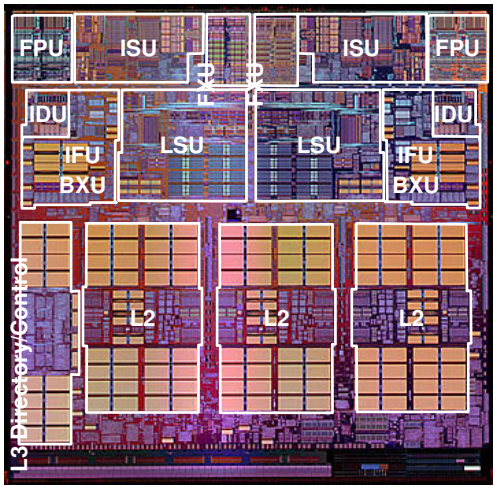

- **First dual-core CPU**
- **On-chip L2 & L3 dir**
- **~1.5 Ghz Frequency**
- **184M transistors**
- **180nm (2001)**
- **130nm (2002) SOI**

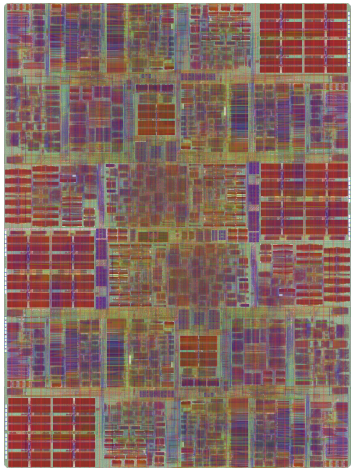

- **Dual core mainframe**
- **~1.2GHz (short pipeline)**
- **121M transistors**
- **130nm (2003) SOI**

IBM Server Characteristics

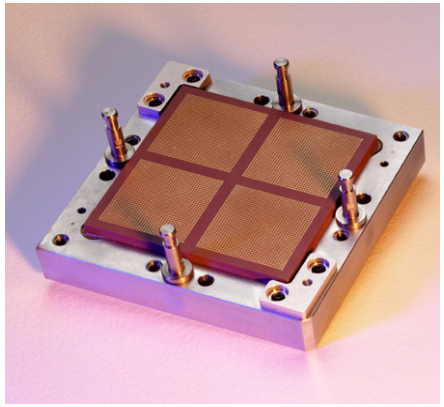

IBM pSeries

- **P690 module**
- **8-way with 4 chips**
- **32-way system**

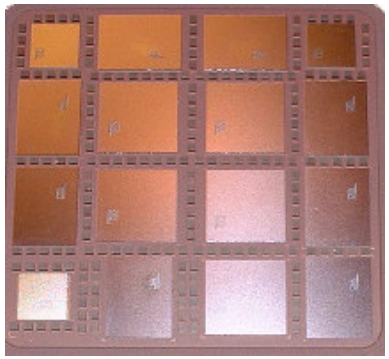

IBM zSeries

- **16-way module**
- **shared L2**

- **Still more parallelism**
- **Will move to chip in future**

- **20 way Freeway 2000**
- **First 64-bit mainframe**
- **64-way Trex 2003**
- **4 16-way boards**

Evolution of Server Design

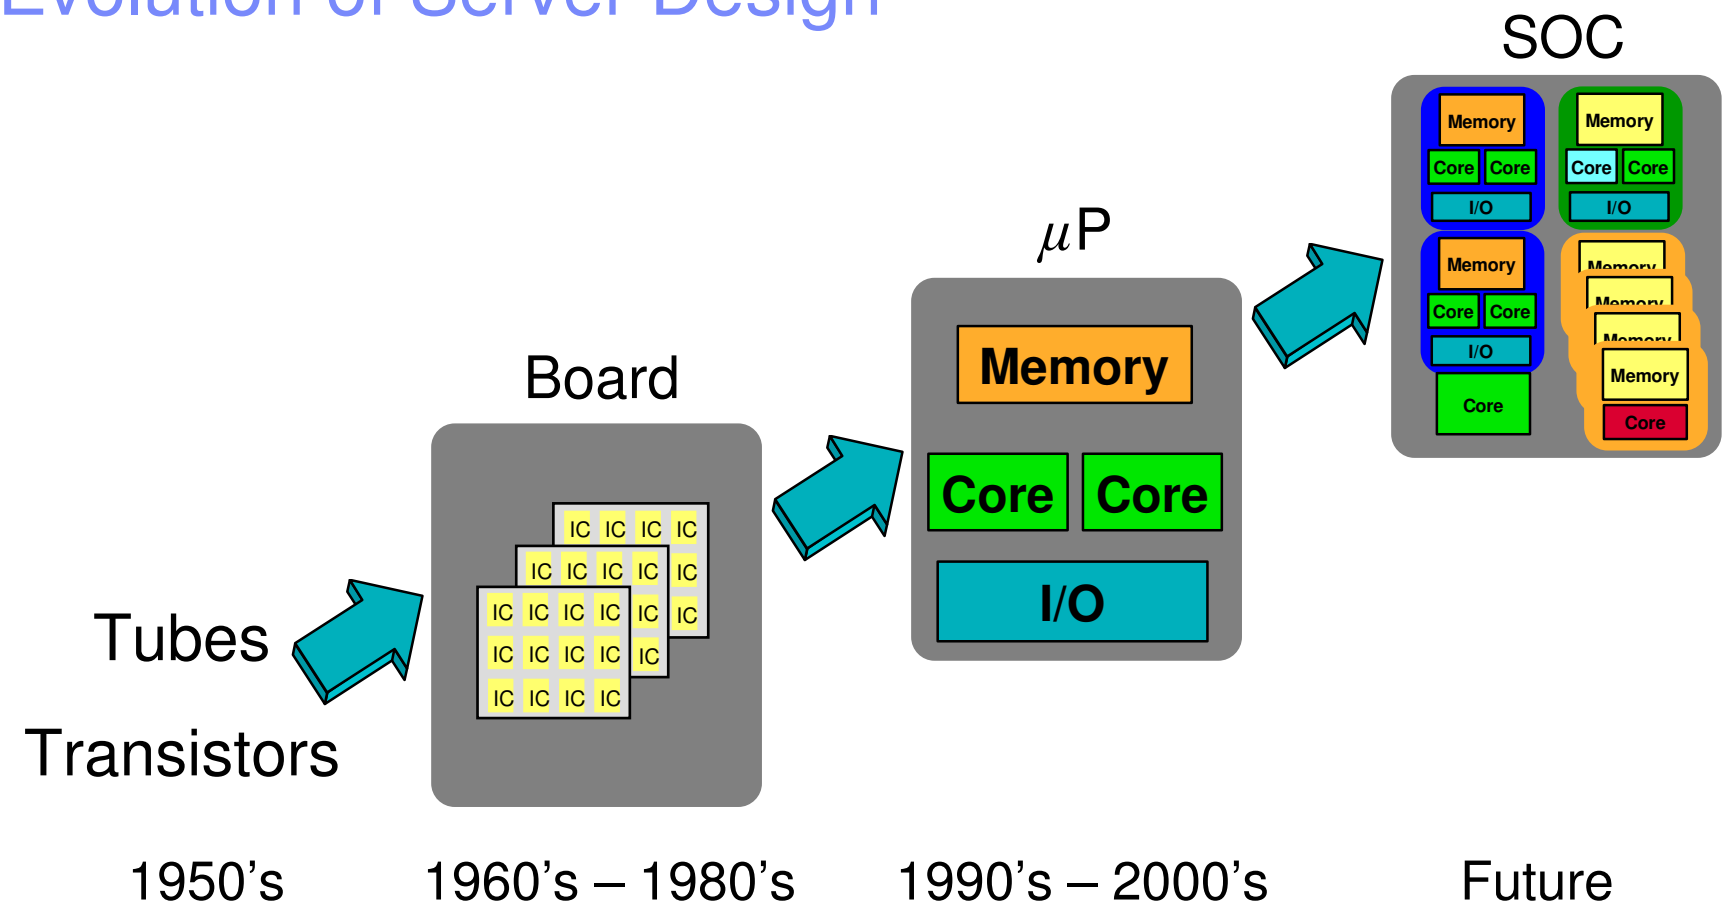

Future Server-on-a-Chip Methodology

- **System-level assembly of components**
- **Efficient custom design of key components**
- **Synthesis of non-custom components**
- **Accurate planning for power, interconnect, timing**
- **Automated chip integration – buffers, clocks, power**

- **Need to combine best of ASIC and processor methods**

Future System-Level Design System

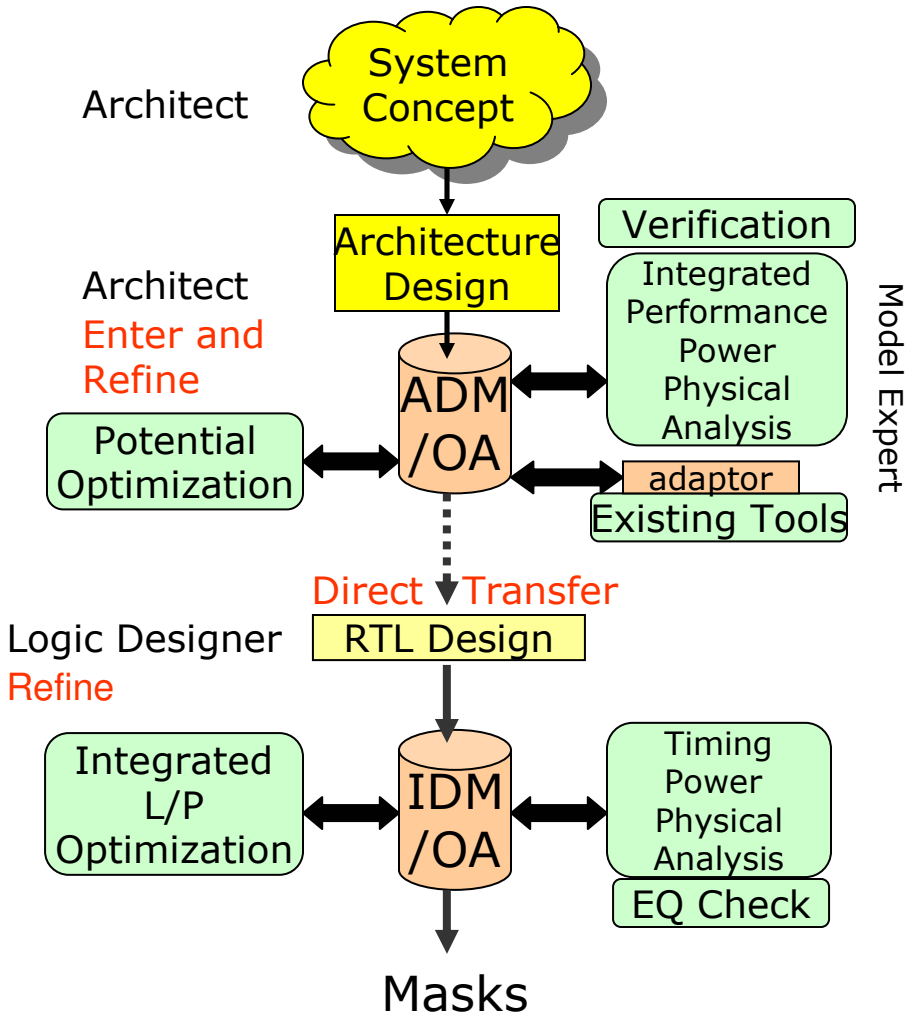

- **System-Level Representation**
- **Integrate power, performance and interconnect analysis**
- **Enable optimization**
- **Link to RTL Implementation**
  - **Accurate prediction**
  - **Accelerate implementation**
- **Verification**
  - **Prove key properties**
  - **Check implementation**
- **Extend Open Access to system level**
  - **Enable tight link to RTL**
  - **Stimulate system-level R&D**

Need System-Level Representation

- **Include all design factors**
  - Function, physical, power, thermal, reliability, …
- **Accept diverse forms**
  - Functional, structural, assertions, constraints, …
- **Support several levels of abstraction**
  - Yet enable verification between levels

System-Level Verification – Data Merger Example

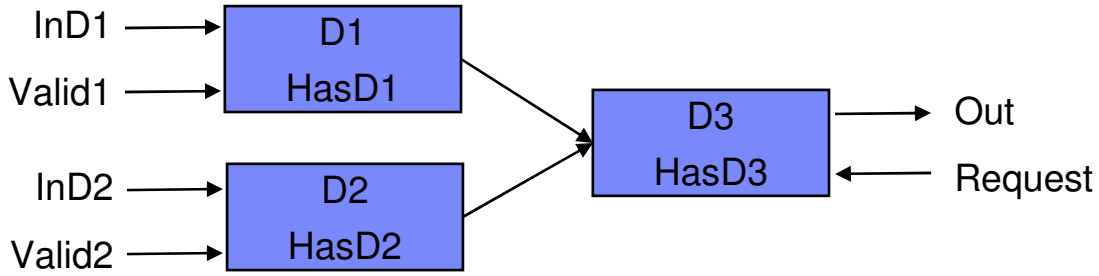

5 Non-deterministic Rules

1. |¬HasD1 & Valid1| → {D1:=InD1; HasD1:=true;}
2. |¬HasD2 & Valid2| → {D2:=InD2; HasD2:=true;}
3. |HasD1 & ¬HasD3| → {D3:=D1; HasD1:=false; HasD3:=true;}
4. |HasD2 & ¬HasD3| → {D3:=D2; HasD2:=false; HasD3:=true;}
5. |Request & HasD3| → {Out:=D3; HasD3:=false;}

System-Level Verification – Experiment

- **Specified Data Merger with Murphi and SystemC**
  - Independent rules simplify task
- **Proved data inputs not lost**
- **Implemented at gate level, VHDL and SystemC**
- **Proved implementation is refinement of specification**
- **Similar results for much larger distributed cache**
- **Greatly reduces proof complexity**

SUMMARY

- **Business demanding more innovative systems**
- **Technology pressuring other performance factors**
- **Systems calling for Server-on-a-Chip methodology**

- **Need to merge best of ASIC and processor methods**
- **Need to raise DA to system-level**
- **Standards can enable faster development**

First Integrated Design System Workshop

- **DAC – Monday, June 13, 2005 Anaheim, California**
- **Follow-on to 5 successful Interoperability Workshops**
- **Organizers**
  - John Darringer – IBM, Rahul Goyal - Intel,
  - Scott Peterson - LSI, Alva Barney – HP
- **Objective**
  - Integrated design systems are essential
  - Open Access is an enabler, but more is required
  - What else is needed for future design systems to keep pace?

Integrated Design System Workshop Agenda

1:00 **Introduction** **John Darringer Research Manager, IBM**

1:20 **User Views** **Chair: Alva Barney, CAD Manager, HP**

Siva Yerramilli, Dir. Design Technology, Intel

David Kung, Sr. Mgr, Design Automation, IBM

Lambert van den Hoven, VP Design Technolgy, Philips

Yoshi Inoue Group Manager Renesas

2:10 **Vendor Views** **Chair: Rahul Goyal, Dir. EDA Business, Intel**

Aurangzeb Khan, VP Design Foundries, Cadence

Eshel Haritan, VP Engineering, Co-Ware

Premal Buch, GM Design Implementation, Magma

Rich Goldman, VP Market Dev, Synopsys

Robert Hum, VP Design, Verification and Test, Mentor Graphics

3:00 **Other Views** **Chair: Scott Peterson, Dir. Rapid Chip Methodology, LSI**

Gary Smith, Chief Analyst, Gartner Dquest

Steve Schulz, CEO, Si2

Andreas Kuehlmann, Dir. Cadence Berkeley Labs

Jim Solomon, CEO Xulu Entertainment

3:45 **Break**

4:10 **PANEL: What's Next?** **Chair: Richard Goering, Editor, EE Times**

Acknowledgements

- **Referenced publications**
  - "Early Analysis Tools for System-on-a-Chip Design", J Darringer, R. Bergamaschi, S. Bhaatacharya, D. Brand, A. Herkersdorf, J. Morrell, I. Nair, P. Sagmeister and Y. Shin, IBM JR&D Vol 46, No 6., Nov. 2002
  - "SEAS: A System for Early Analysis of SoCs", R. Bergamaschi, Y. Shin, N. Dhanwada, S.Bhattacharya, W. Dougherty, I. Nair, J. Darringer, S. Paliwal, CODES-ISSS 2003
  - "Formal Verification of Distributed Cache Memory Protocols", S. German, FMCAD 2004
  - "Understanding Basic Tradeoffs in Multiple Frequency Processor Core Micro-architectures through a Statistical Timing Model", I. Nair, A. Buyuktosunoglu, internal paper
  - "An efficient System-on-a-Chip Design methodology for Networking applications", V. Salapura, C. Georgiou and I. Nair, CASES '04
- **Plus discussions with**
  - Jeff Welser, Erik Kronstadt, Jaime Moreno,