

Hot Interconnects'06 Panel #2

Challenges for Future Interconnection Networks: Power, Reliability, Performance Scalability?

Panel Moderator:

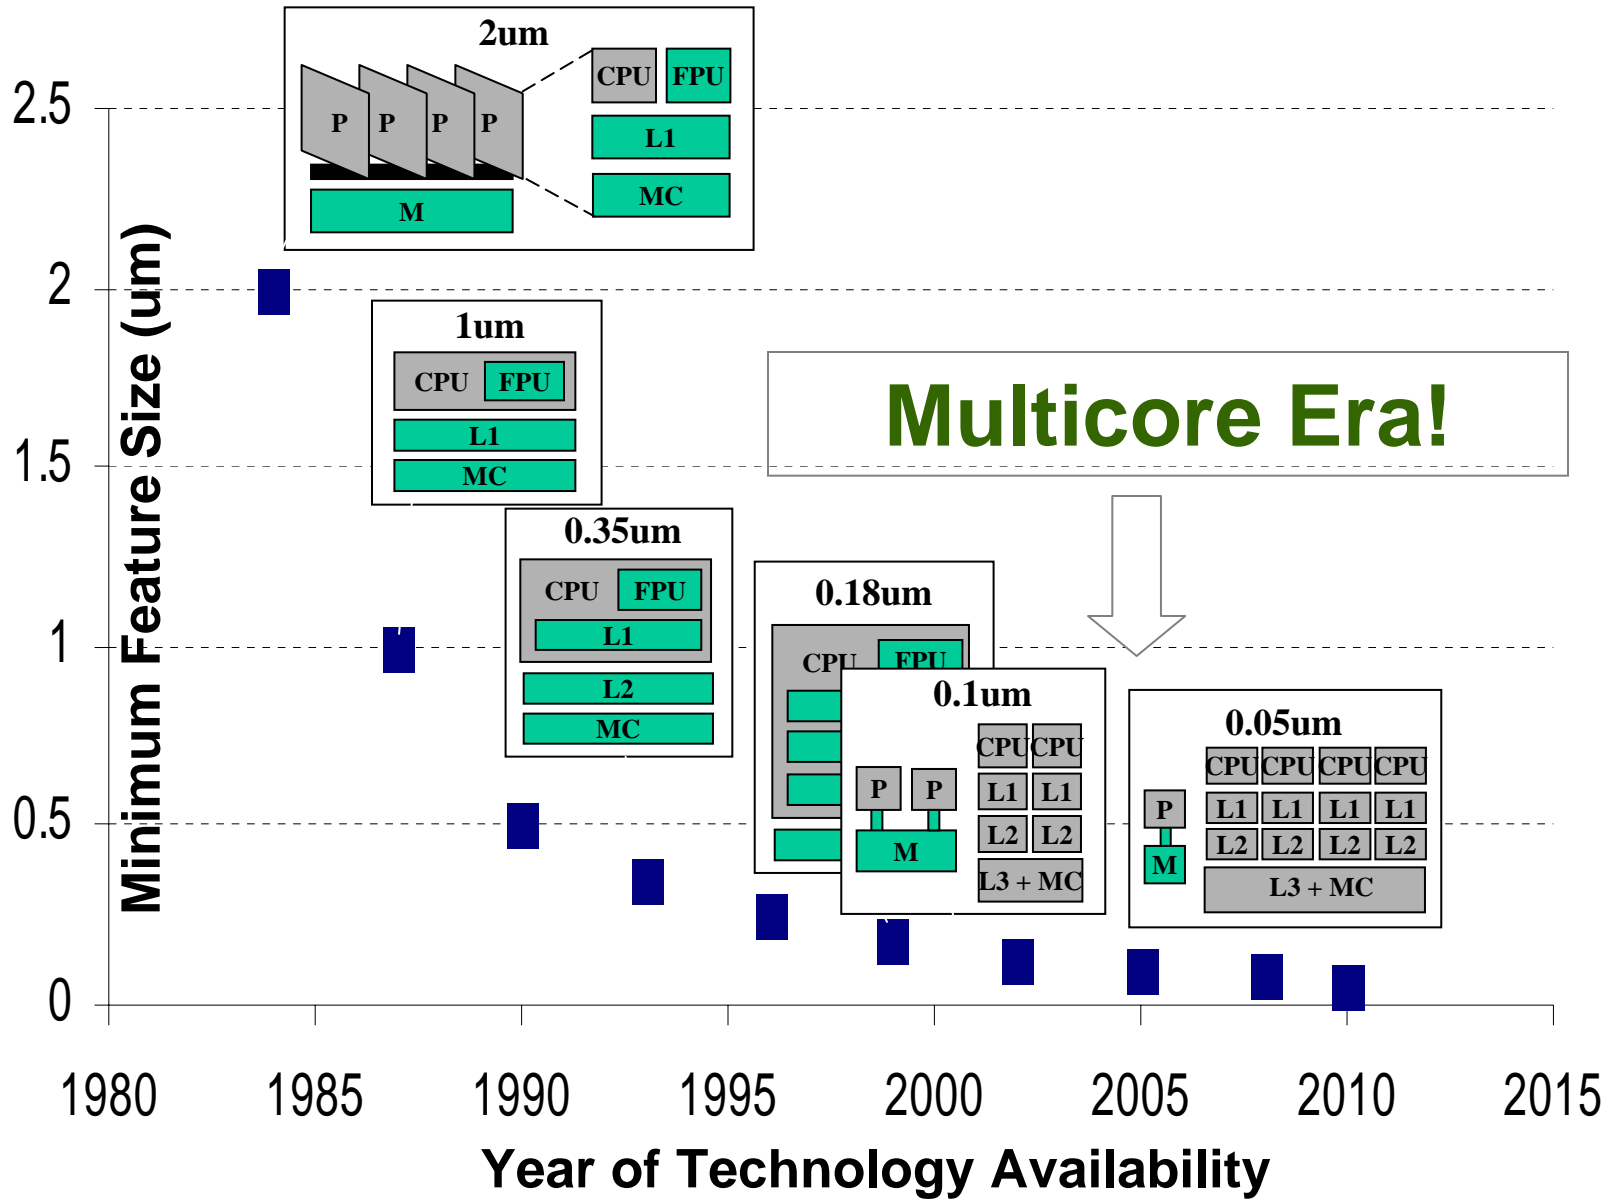
Timothy M. Pinkston

National Science Foundation (NSF)

tpinksto@nsf.gov

University of Southern California (USC)

Proliferation of Interconnects




Panelists

- ◆ Dennis Abts, Cray
- ◆ Bill Dally, Stanford University
- ◆ Jay Jayasimha, Intel
- ◆ Doug Joseph, IBM T. J. Watson Research Labs
- ◆ Steve Keckler, University of Texas at Austin
- ◆ Dhableswar (D. K.) Panda, The Ohio State University

Discussion Questions

- ◆ Ultimately, what will determine latency and throughput performance: the interfaces to the network or the core network fabric?

- ◇ The network can be considered as a “*pipe*” of variable width

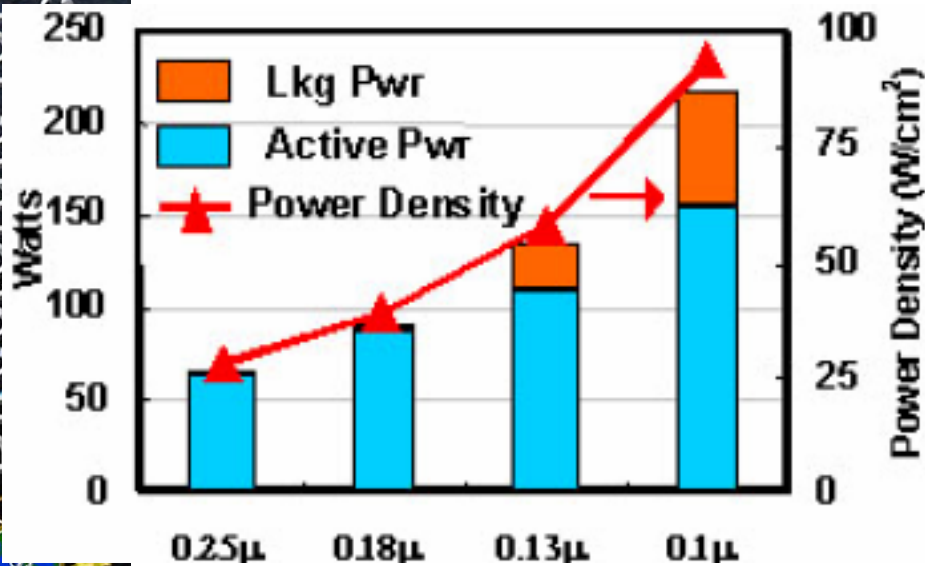


- ◇ There are three points of interest *end-to-end*:
 - ◆ Injection into the pipe
 - ◆ Narrowest section within pipe (i.e., min. bisection having traffic)
 - ◆ Reception from the pipe
- ◇ ***The bandwidth at the narrowest point and utilization of that bandwidth determines the throughput!!***

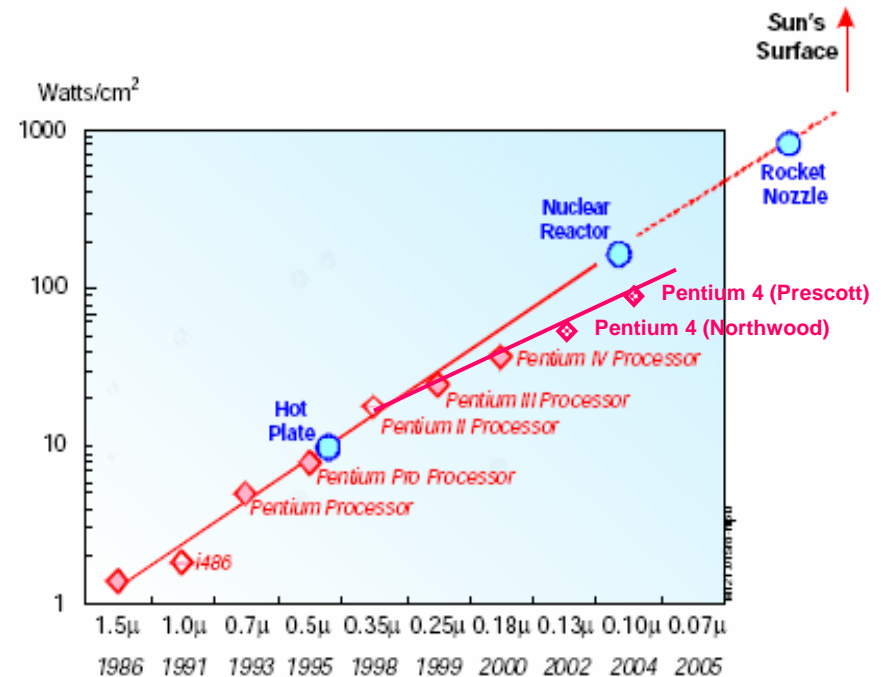
“Interconnection Networks”, Appendix E in *Computer Architecture: A Quantitative Approach*, 4th Edition, Elsevier Publishers, 2006.

Discussion Questions

- ◆ Will power consumption mainly be the problem of the processor and memory system designers or will it mainly be the network designer's problem?
 - Power wall being reached with voltage scaling limitations, thermal problems, operational variation, etc.



Source: S.Ruso, ESSCIRC 2001



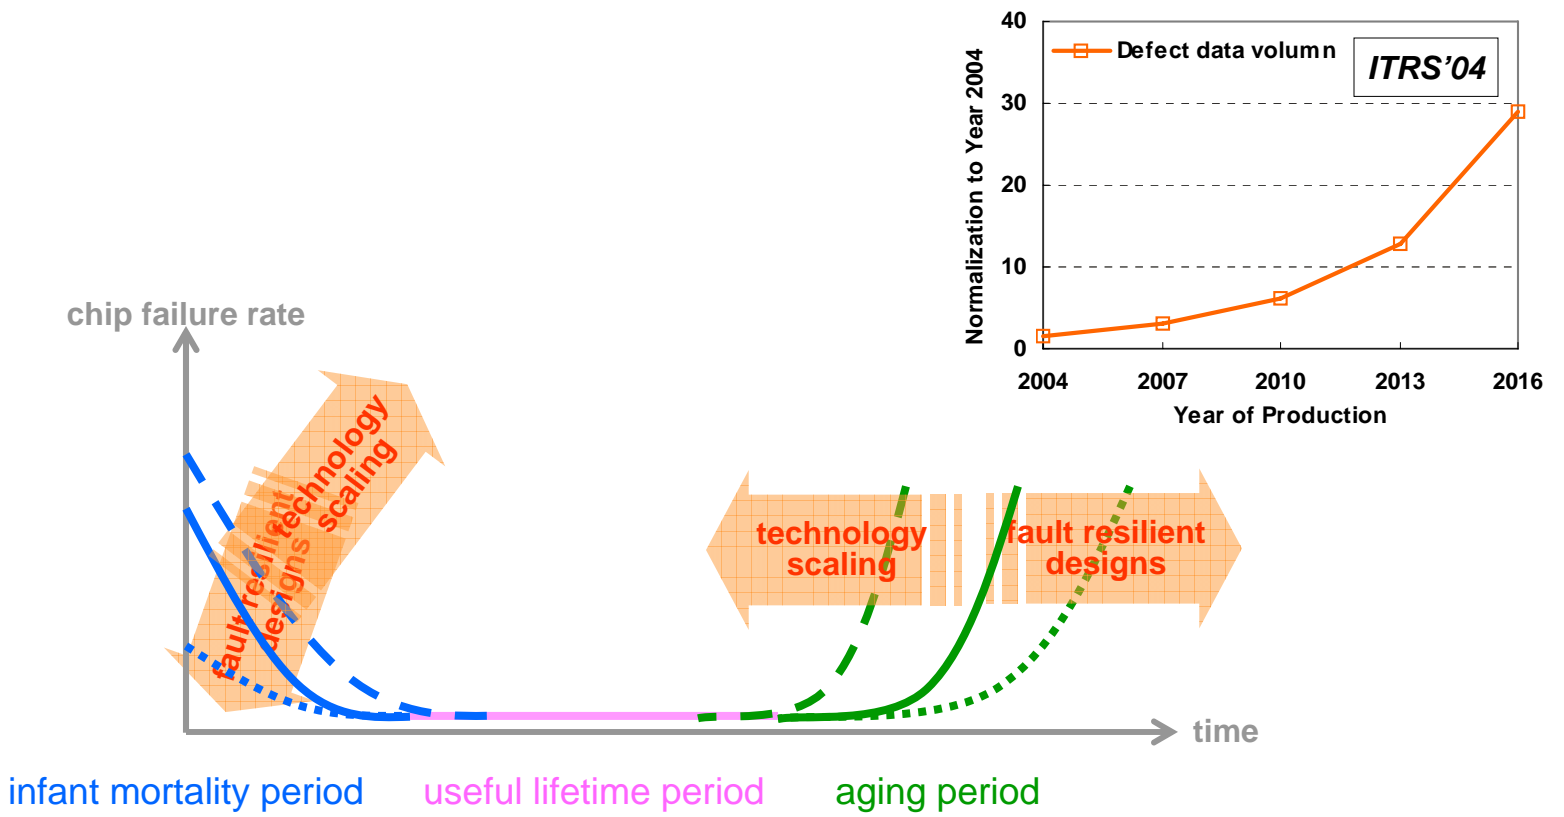
Source: Prismark Partners LLC, Feb 2002

Discussion Questions

- ◆ What's wrong with simply turning down the clock frequency and going serial as a means of confronting the power problem in networks?
 - Blue Gene/L
 - PCI-Express
- ◆ What critical reliability assumptions we now take for granted are likely not to hold for future interconnection networks, and what might be the impact?
 - low MTTF and MTTR
 - field replaceable
 - hot swappable

Discussion Questions

- ◆ How drastically do network reliability concerns grow as we move to nanoscale and multi-billion device integration densities on chip?



Discussion Questions

- ◆ Are there any new, promising technologies and/or design/architecture paradigms on the horizon that may swoop down and save the day?
 - quantum, nano, 3D integration, optical, reconfigurable, self-repairable/adaptable, multicore architectures, etc.
- ◆ On-chip networks—the wave of the future? Take a pick:
 - same problems, same solutions?
 - same problems, different solutions?
 - different problems, same solutions?
 - different problems, different solutions?

Panelists

◆ Dennis Abts, Cray

◆ Bill Dally, Stanford University

◆ Jay Jayasimha, Intel

◆ Doug Joseph, IBM T. J. Watson Research Labs

◆ Steve Keckler, University of Texas at Austin, USA

◆ Dhabalaswar (D. K.) Panda, The Ohio State University



NSF 2006 Proposal Solicitations

CCF Division's CPA 2006 Solicitation 06-585

- ◇ *Proposal due date: October 10, 2006*
- ◇ *<http://nsf.gov/pubs/2006/nsf06585/nsf06585.htm>*

CNS Division's CSR 2006 Solicitation

- ◇ *Proposal due date: November 10, 2006*
- ◇ *<http://nsf.gov/pubs/2005/nsf05629/nsf05629.htm>*

◆ Research topics include

- ◇ **SW/HW systems:** reliable and high-performance computing, parallelizing compilers, programming models, and run-time support for efficient resource allocation and scheduling
- ◇ **Computer system architecture:** processor μ architecture, memory, I/O subsystems, interconnection networks (including on-chip networks), reconfigurable and application-specific architectures; multicore, multithreaded, and SoC architectures
- ◇ **SW/HW tools:** design, simulation, benchmarking, performance measurement, evaluation and tuning