

---

# Challenges for Future Interconnection Networks

William J. Dally  
Computer Systems Laboratory  
Stanford University

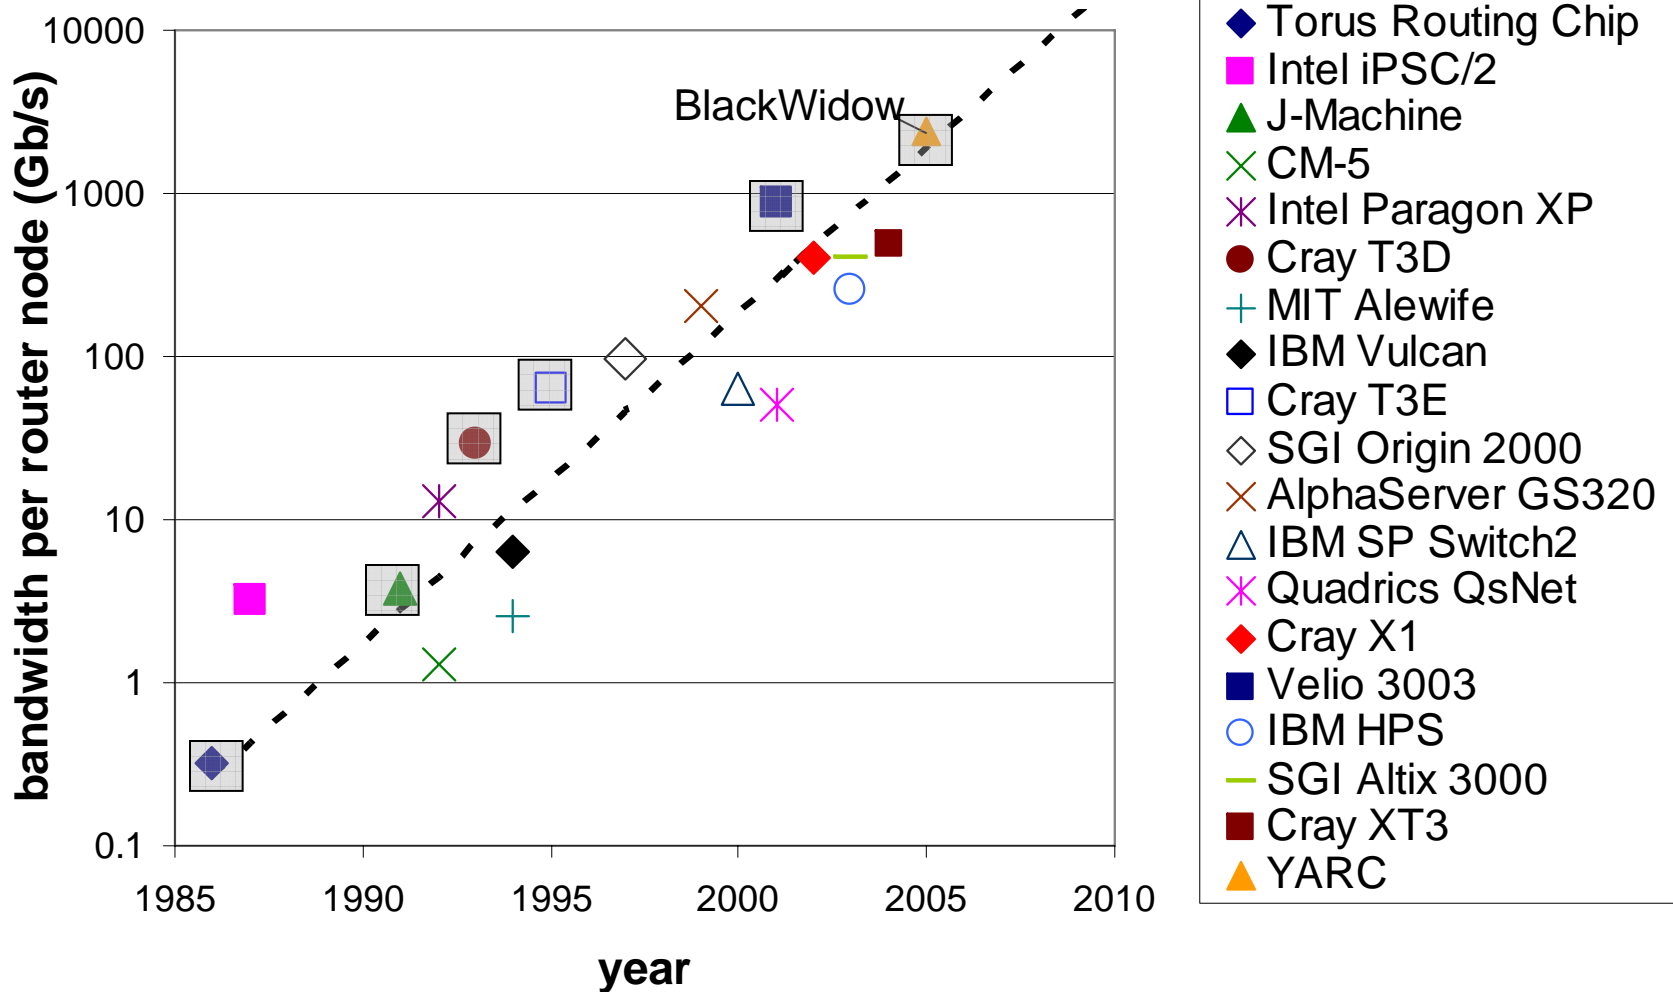
Hot Interconnects Panel  
August 24, 2006

# Interconnection Networks are Critical

---


- Interconnection networks are becoming the dominant component of future systems
  - Bulk of power is in data movement
  - Latency of critical paths dominated by interconnect
  - Bandwidth limited by interconnect
  - Arithmetic is (almost) free - data/instruction movement is what matters
- But today most interconnects are ad-hoc
  - Small number of terminals (2-8)
  - Connected with bus, crossbar, ring
  - Little analysis or optimization - until bottleneck discovered

# Technology Trends...



# Today - We can build very capable system-wide networks

---




# Many Recent Advances

---

- High-radix networks
- Global adaptive routing
- Tools for topology optimization

To find out more about networks, read




Hot-I: 6


Aug 24, 2006

# But much of the future is on-chip (CMP, SoC, Operand)


---




2006



2007.5




2009




2010.5

Hot-I: 7



2012

2013.5


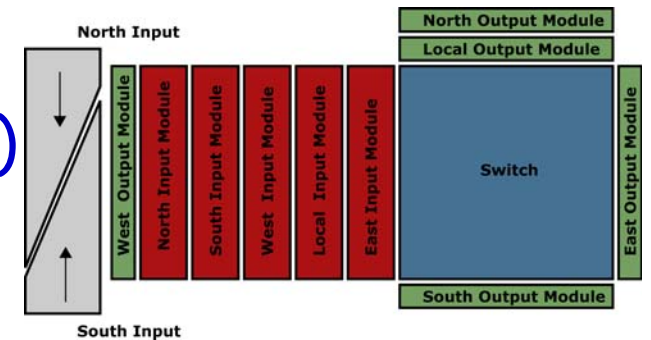


2015

Aug 24, 2006

# On-Chip Networks are Fundamentally Different


- Different cost model
  - Wires plentiful, no pin constraints
  - Buffers expensive (consume die area)
  - Slow signal propagation
- Different usage patterns
  - Particularly for SoCs
    - Significant isochronous traffic
    - Hard RT constraints
- Different design problems
  - Floorplans
  - Energy-efficient transmission circuits



# Large room for improvement over 'obvious' approaches

---

Network Energy  $\times$  Completion Time  
(normalized to Torus network)



Source: Balfour et al. ICS-06

# Challenges and Directions

---

- Prototype to get real data on NoC cost and performance
- Topology and flow-control to exploit plentiful bandwidth to save buffers and to support real-time and isochronous traffic
- System-level design of distributed processor/storage hierarchy/network (NIC overhead often dominates)
- Power and reliability not overwhelming issues - but we need to think of a network as a communication system  $E_b/N_0$
- Get people to treat interconnect design as a first-class endeavor (not an afterthought)

# Summary/Rebuttal

---

- DK - Software needed to manage multi-core systems (data placement)
- DA - High radix is good
- JJ - Learn to live in "flatland" where network provides fault tolerance
- DJ - Link-level retry is good (RR).  
Strong scaling = fine granularity
- SK - Real on-chip networks
- BD - On chip networks are a new frontier - we need to get our hands dirty with them