Slide 0

Hardware Support for Controlled Interaction of Guaranteed and Best-Effort Communication

Jennifer Rexford, James Dolter, and Kang Shin

April 28, 1994

Real-Time Computing Laboratory
The University of Michigan
Ann Arbor, Michigan

Slide 1

Motivation

- Point-to-point networks
  → fault-tolerance, locality, scalability
- Hardware communication support
  → relieve software, network utilization, latency
- Controlled sharing of network resources
  → traffic classes, performance metrics, cost

Slide 2

Routing and Switching

Routing

- Static vs. adaptive
- Minimal vs. nonminimal

Switching

- Packet switching
- Virtual cut-through
- Wormhole

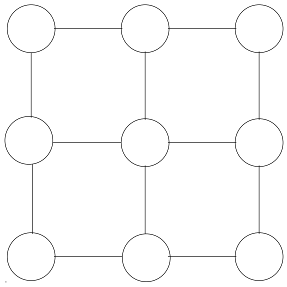

Real-Time Computing Laboratory
The University of Michigan

Slide 3

Switching Schemes and Resource Utilization

Switching scheme	**Memory/Processing**	**Link/Channel**
Packet switching	Predictable	Predictable
Virtual cut-through	Unpredictable	Predictable
Wormhole	Predictable	Unpredictable

Real-Time Computing Laboratory
The University of Michigan

Slide 4

Traffic Classes

Guaranteed: Strict performance bounds

- Schedule packet transmissions
- Reserve bandwidth

→ packet switching and static routing

Best-effort: Average performance

- Avoid packet buffering
- Bypass congestion

→ cut-through switching and adaptive routing

Real-Time Computing Laboratory
The University of Michigan

Slide 5

Scalable Point-to-point Interface DrivER (SPIDER)

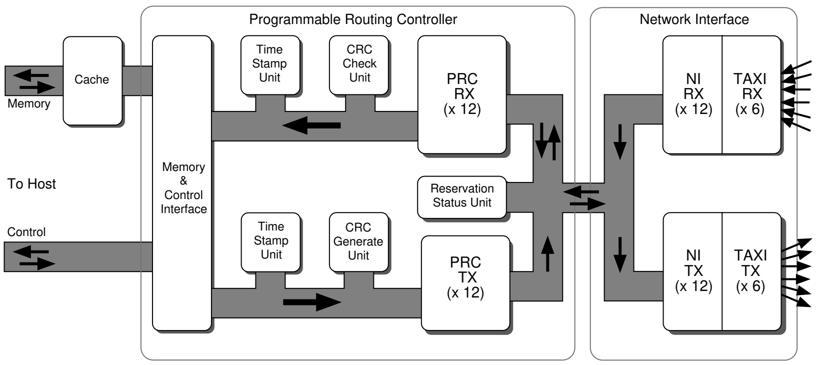

Real-Time Computing Laboratory
The University of Michigan

Slide 6

Intrusion of Best-Effort Traffic

Best-effort packets affect network access time for guaranteed packets:

- Single FIFO for all packets → service time for queue
- Separate queues for traffic classes → service time for one packet

This limits:

- Pipelining of best-effort packets
- Packet size for best-effort messages
- Usage of wormhole switching for best-effort packets
- Admission criterion of guaranteed traffic

Real-Time Computing Laboratory
The University of Michigan

Slide 7

Partitioning Traffic Classes

- Multiple virtual channels per link
- Partition traffic onto virtual networks
- Tailor low-level policies for each virtual network
- Channel contention in one class does not harm other traffic
- Fine-grain arbitration amongst virtual channels

→ Controlled, fine-grain sharing of link bandwidth

Real-Time Computing Laboratory
The University of Michigan

Slide 8

Controlled Interaction of Traffic Classes

- Cycle-level simulation model of SPIDER
- 6×6 torus (wrapped square mesh)
- Dimension-ordered static routing
- Virtual networks
  - Best-effort: wormhole switching
  - Guaranteed: packet switching
- Fair arbitration amongst virtual channels

Real-Time Computing Laboratory
The University of Michigan

Slide 9

Average Latency

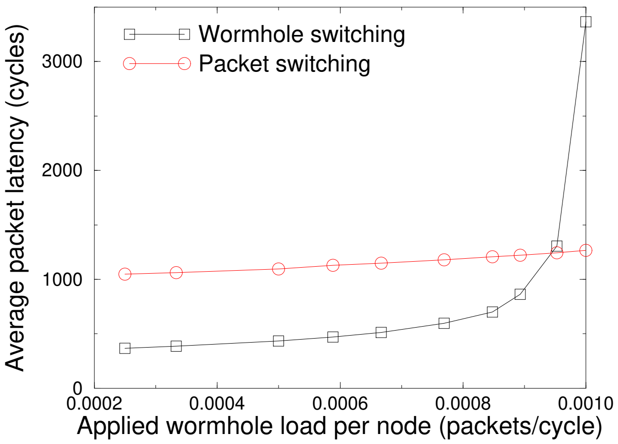

Real-Time Computing Laboratory
The University of Michigan

Slide 10

Latency Standard Deviation

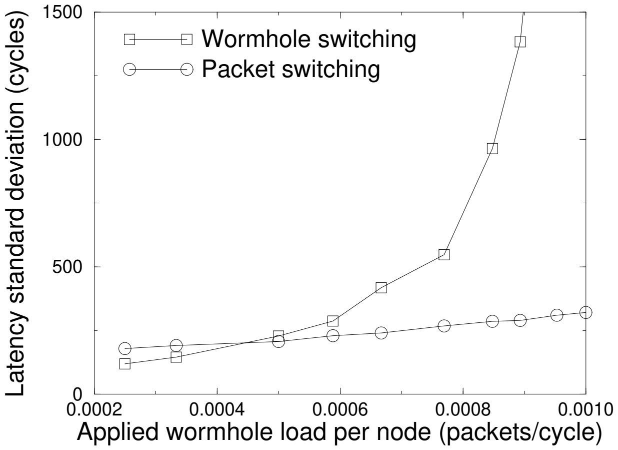

Real-Time Computing Laboratory
The University of Michigan

Slide 11

Conclusions and Future Work

- Each routing-switching scheme best-suited to certain traffic
- Hardware support for multiple low-level policies
- Virtual networks for traffic partitioning
- Fine-grain arbitration for controlled interaction between classes
- Interaction of routing, switching, queueing, and arbitration

→ best-effort traffic can utilize large packets, cut-through switching, and adaptive routing without restricting the admission of guaranteed traffic

Real-Time Computing Laboratory
The University of Michigan