Slide 0

Support for Multiple Classes of Traffic in Multicomputer Routers

Jennifer Rexford and Kang G. Shin

May 17, 1994

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

Slide 1

Outline

- Related work and motivation
- Guaranteed and best-effort traffic
- Evaluation of switching schemes
- Traffic mixing
- Conclusions and future work

Slide 2

Related Work on Traffic Mixing

- Control and data communication [Dally87, Dally92]
- Coarse and fine-grain communication [Peterson91]
- Long and short messages [Kim93, Konstantinidou94]
- Real-time scheduling in wormhole networks [Li94, Mutka94]

Slide 3

Real-Time and Multimedia Applications

- **Applications:** scientific visualization, process control, instrumentation
- **Performance metrics:** average latency, latency jitter, throughput
- **Traffic classes:** guaranteed and best-effort

→ application performance requirements affect the suitability of architectural features in multicomputer routers

Slide 4

Traffic Classes

Guaranteed: Strict performance bounds

- Schedule packet transmissions
- Reserve bandwidth
- Bound network access time

Best-effort: Good average performance

- Avoid packet buffering
- Bypass congestion

Slide 5

Architectural Features

Architectural parameters affect resource usage by packets

- **Routing:** where node/link resources are used
- **Switching:** what type of buffer/channel resources are used
- **Queueing:** order/scheduling of resource acquisition
- **Arbitration:** interleaving access amongst competing packets

Slide 6

Conflicting Goals for Guaranteed and Best-Effort Traffic

	Guaranteed Traffic	**Best-Effort Traffic**
Switching	Packet switching	Wormhole switching
Routing	Static	Adaptive
Queueing	Centralized	Decentralized
Flow control	Packet	Flit
Packet size	Fixed	Variable

Slide 7

Switching Schemes and Latency

Switching schemes: wormhole, virtual cut-through, and packet switching

Latency metrics: average and coefficient of variation

Flit-level simulations:

- 8×8 torus (8-ary 2-cube)
- Poisson generation of 16-flit packets at each node
- Uniform random selection of destination nodes

Real-Time Computing Laboratory
The University of Michigan

Slide 8

Average Packet Latency

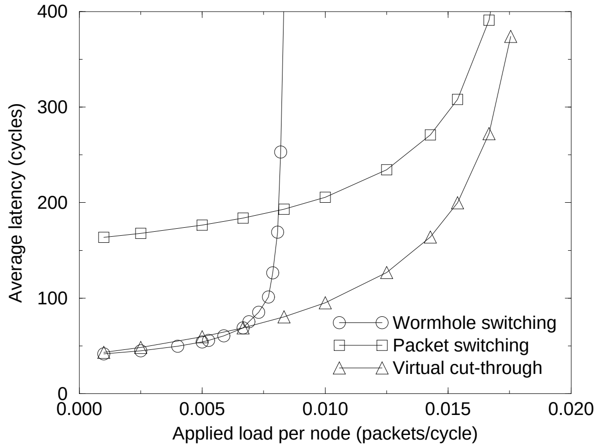

Real-Time Computing Laboratory
The University of Michigan

Slide 9

Latency Coefficient of Variation (5-hop packets)

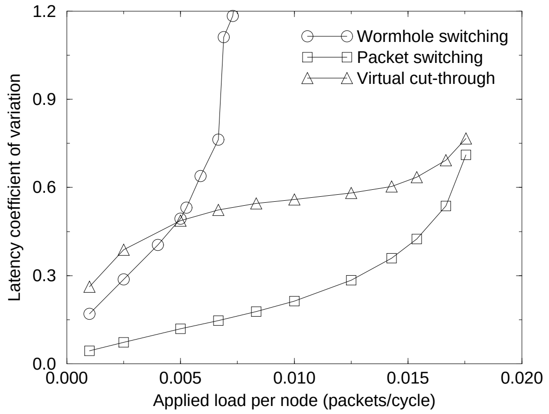

Real-Time Computing Laboratory
The University of Michigan

Slide 10

Wormhole Switching

Queueing: Order of virtual channel reservation

Arbitration: Interleaving of virtual channels at physical link

- Queueing decentralized across multiple nodes
- Stalled wormhole packet dilates its length
- Demand-driven arbitration yields variable flit service rate
- Link sits idle when all virtual channels are blocked

Slide 11

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

Slide 12

Coexistence of Guaranteed and Best-Effort Traffic

- Multiple virtual channels per link
- Separate virtual network for each traffic class
- 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

Slide 13

Evaluation of Traffic Mixing

Guaranteed load: Different fixed guaranteed (packet switching) loads

Latency metrics: Average latency for best-effort and guaranteed packets

- Virtual networks
  - Best-effort: wormhole switching (2 virtual channels)
  - Guaranteed: packet switching (1 virtual channel)
- Dimension-ordered static routing
- Demand-slotted, round-robin arbitration amongst virtual channels
- Vary best-effort (wormhole) load

Slide 14

Average Best-Effort Latency

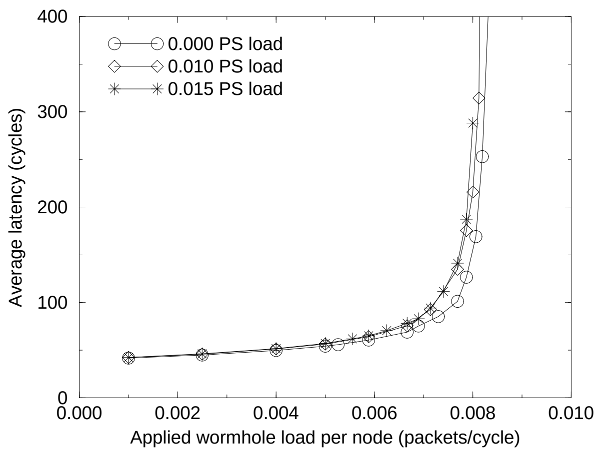

Real-Time Computing Laboratory
The University of Michigan

Slide 15

Average Guaranteed Latency

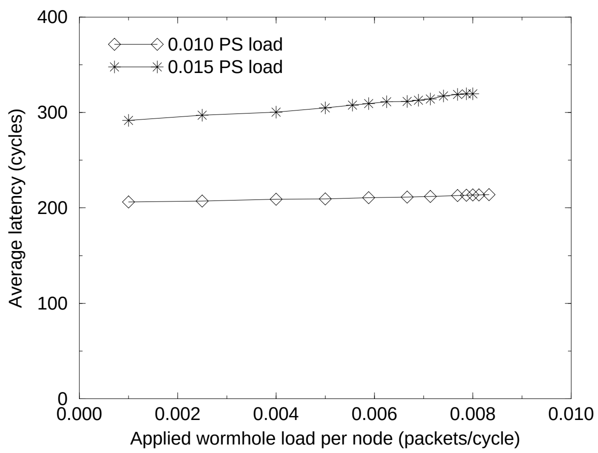

Real-Time Computing Laboratory
The University of Michigan

Slide 16

Tightening Bounds for Guaranteed Traffic

- Priority arbitration to favor guaranteed packets
  - Flit-level preemption of best-effort packets
  - Abstractions for scheduling/allocation for guaranteed traffic
  - Poor performance for best-effort traffic as load increases
- Credit scheme for best-effort packets
  - Up to α best-effort flits can accompany guaranteed packets
  - Bounded dilation of length ℓ guaranteed packet to $\ell + \alpha$
  - Relieve contention in the best-effort virtual network

Real-Time Computing Laboratory
The University of Michigan

Slide 17

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
- Abstractions for bandwidth allocation and packet scheduling

→ 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