Gigabit Routing and Quality of Service

Networld+Interop Session C22

Stephen Haddock

Extreme Networks

Myths of Local Area Networking

- **Myth #1: You cannot route at gigabit speeds.**
  - Preconceived belief that router performance will always lag switching performance.
- **Myth #2: You cannot get Quality of Service on Ethernet networks**
  - Preconceived belief that QoS can be provided only on connection oriented networks with fixed size transmission units.

Routing at Wire Speed

- **First separate the routing functions into protocol processing and packet processing**
- **Protocol Processing:**
  - e.g. RIP, OSPF, DVMRP, IGMP, RSVP
  - complex (many operations, many decision points)
  - event driven, infrequent operation
  - required processing power scales with network size
  - ideally suited to implementation with software modules running on a general purpose CPU

Routing at Wire Speed (cont.)

- **Packet Processing:**
  - e.g. next hop determination, new MAC DA, TTL decrement, checksum update, MAC SA insertion
  - simple, well-defined operations
  - identical operations performed on every packet
  - required processing power scales with network bandwidth, packet rate
  - ideally suited to dedicated hardware processing elements

How to Route at Wire Speed

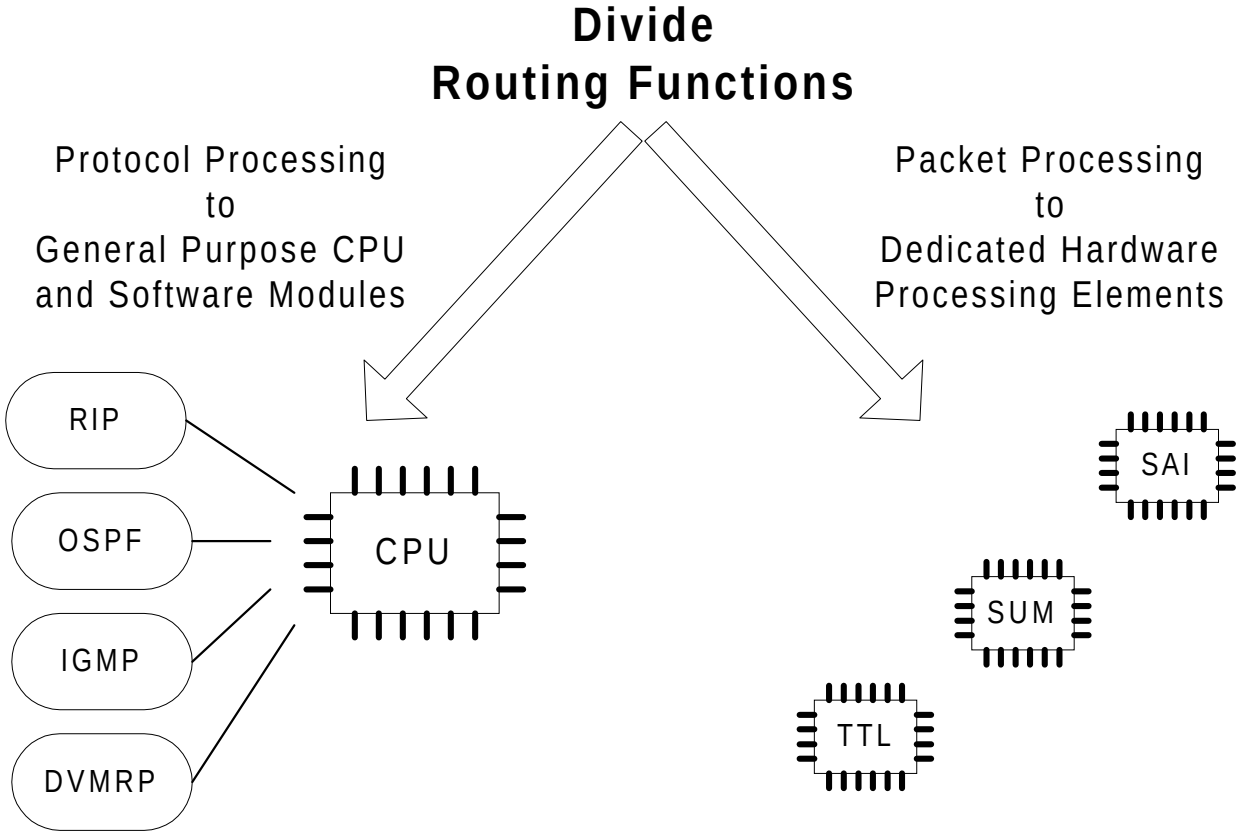

Pipeline the Processing Elements

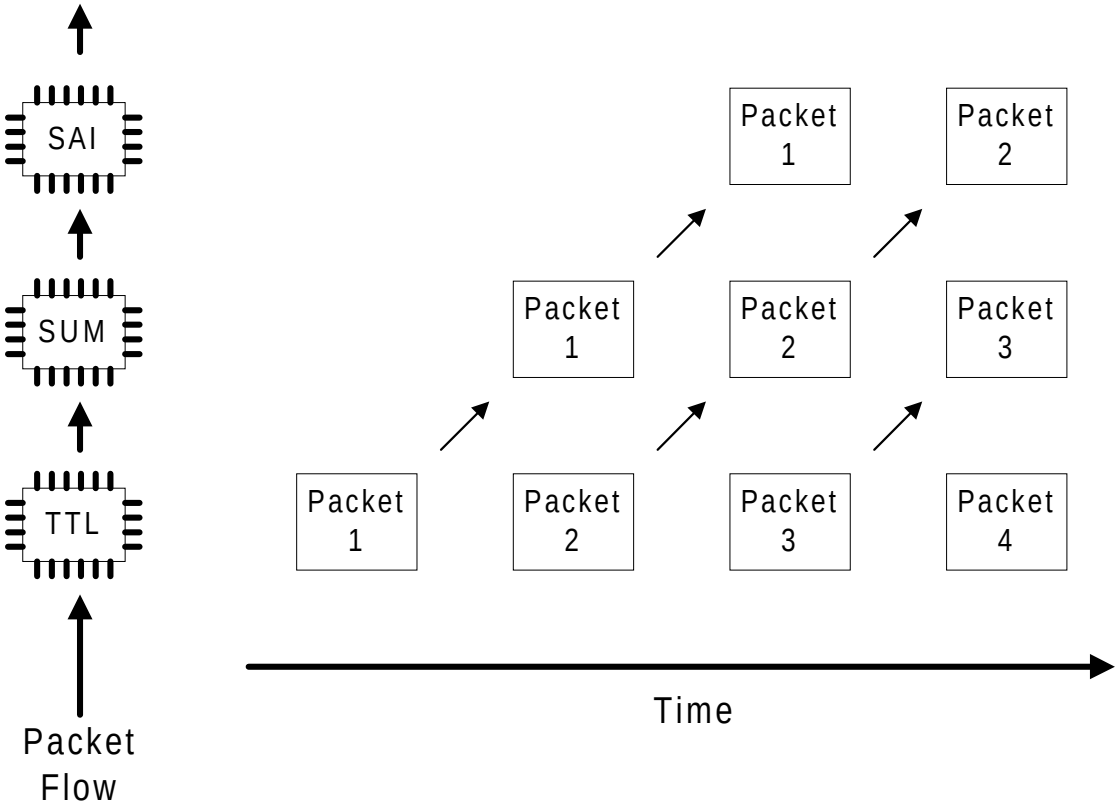

Use Parallel Processing Pipelines

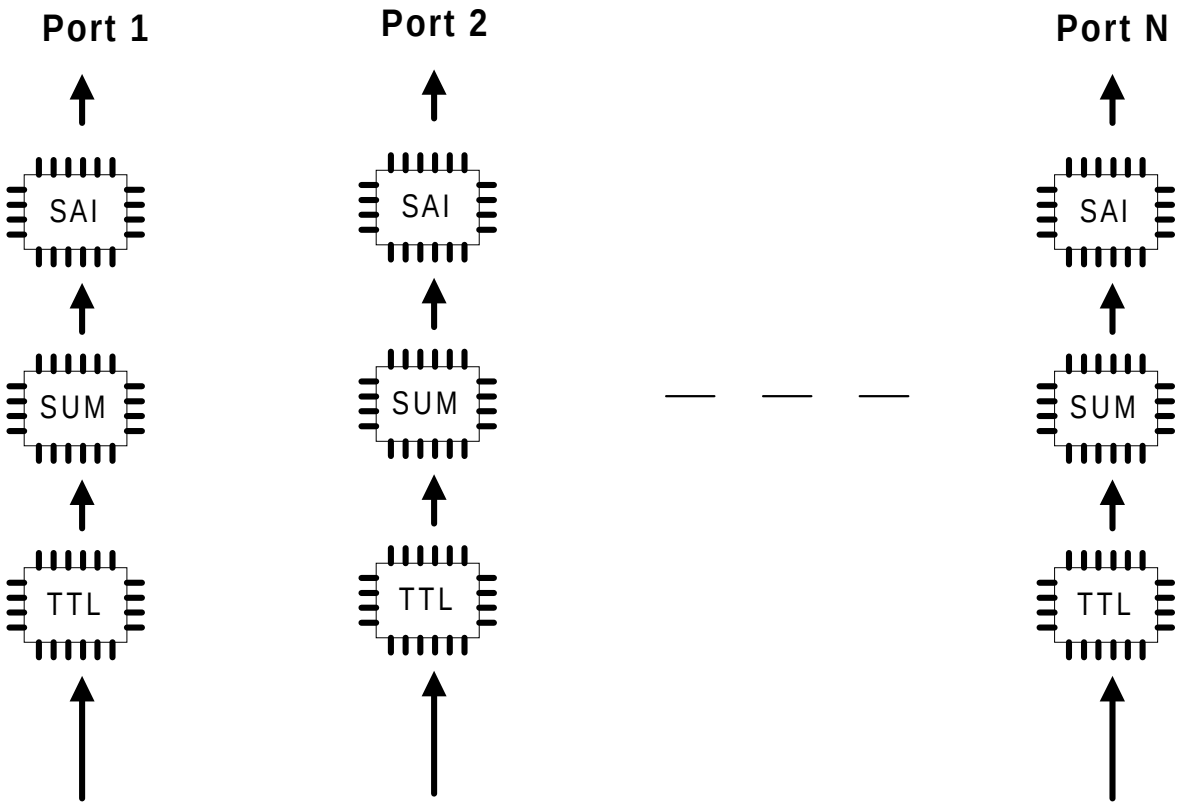

Add the Switch Fabric

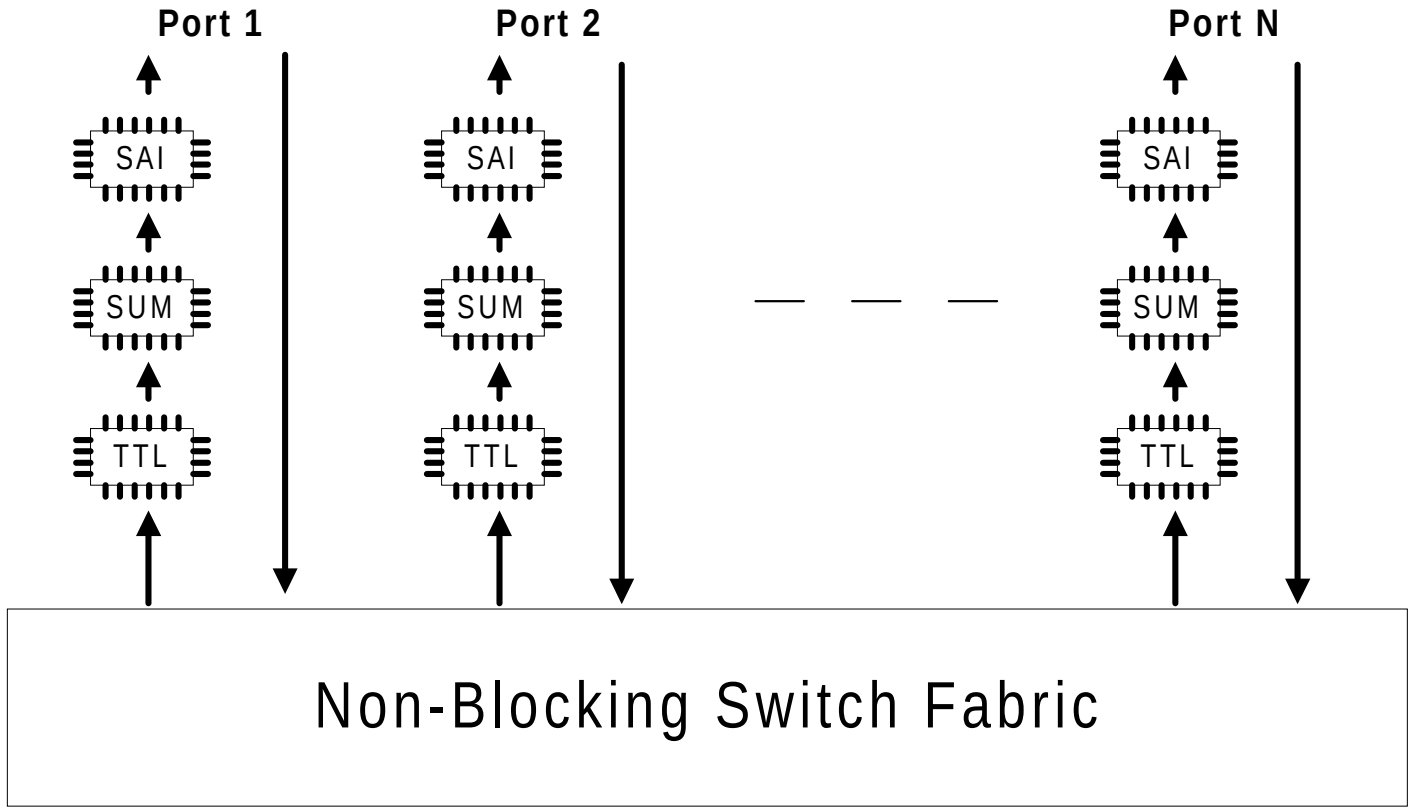

Choice of Switch Fabric

- **Central shared memory chosen as switch fabric in order to meet QoS objectives**
  - Flexible allocation of buffer resources
  - Flexible queueing enables QoS guarantees
  - Efficient multicast handling
  - Efficient CPU access to any packet
  - Most cost effective memory structure
- **But very high bandwidth requirements**
  - Key is a scaleable memory structure

Real Time Forwarding Decision

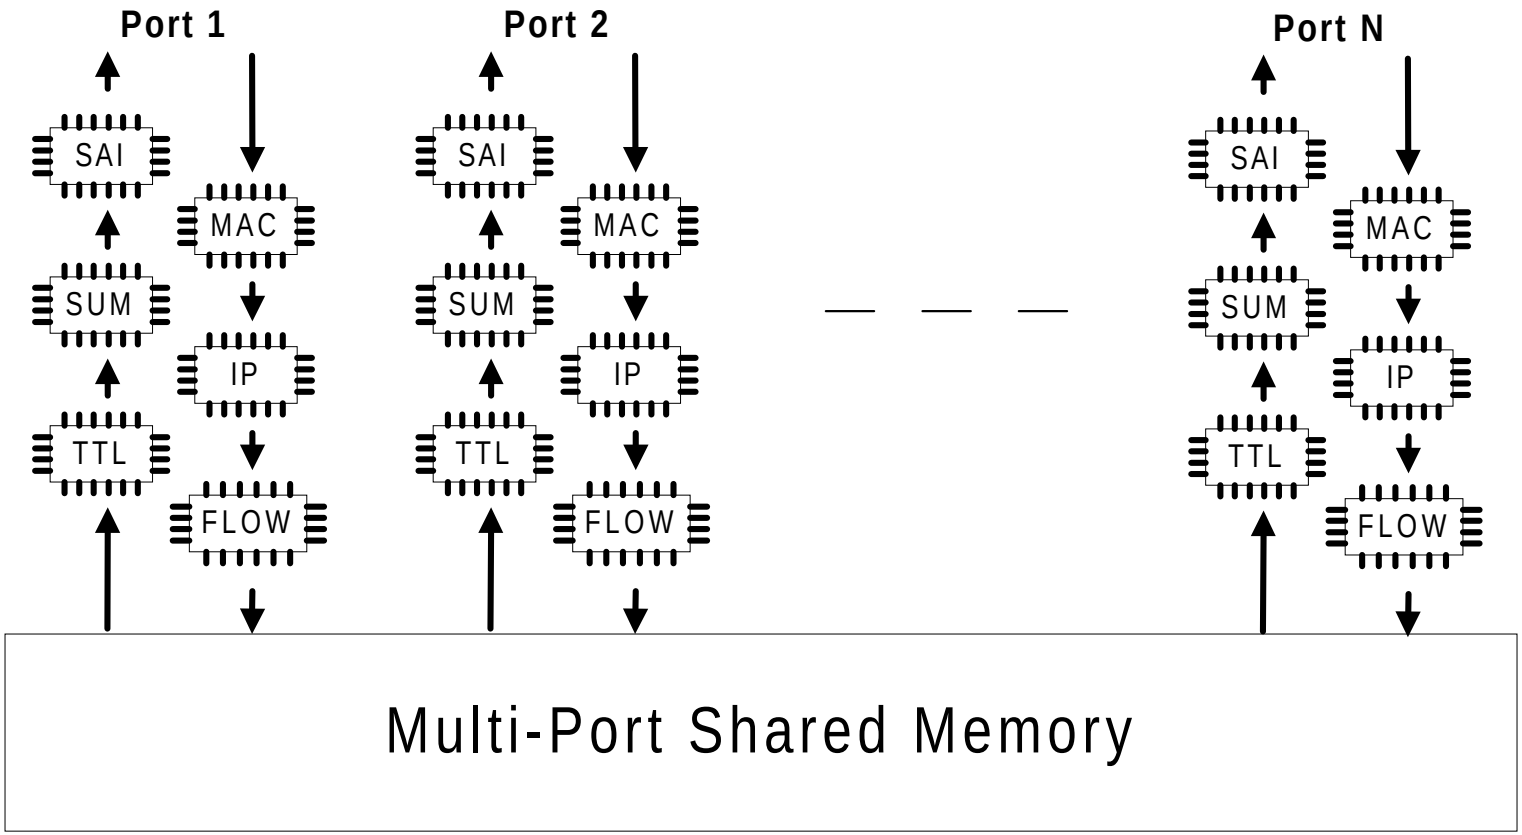

Making the Forwarding Decision

- **Decision process is a sequence of comparisons against a Forwarding Data Base**
  - Simple, repetitive operations well suited for parallel pipelined processing
- **Central vs. Distributed Data Base**
  - Centralized is most efficient to update and most cost effective, but requires very high bandwidth
- **Same memory architecture that enables centralized packet buffering, also enables centralized Forwarding Data Base**

Quality of Service on Ethernet

- **Evolution of Ethernet**
  - Shared Ethernet segments connected by routers
  - Ethernet collision domains separated by layer-2 switches connected to a routed backbone
  - Point-to-point full duplex Ethernet connects layer-3 switches at the core of the network
- **Queuing structures allow differential service levels even on shared links**
- **Switched network of point-to-point full duplex links enables delivery of guaranteed services**

QoS Switch Model

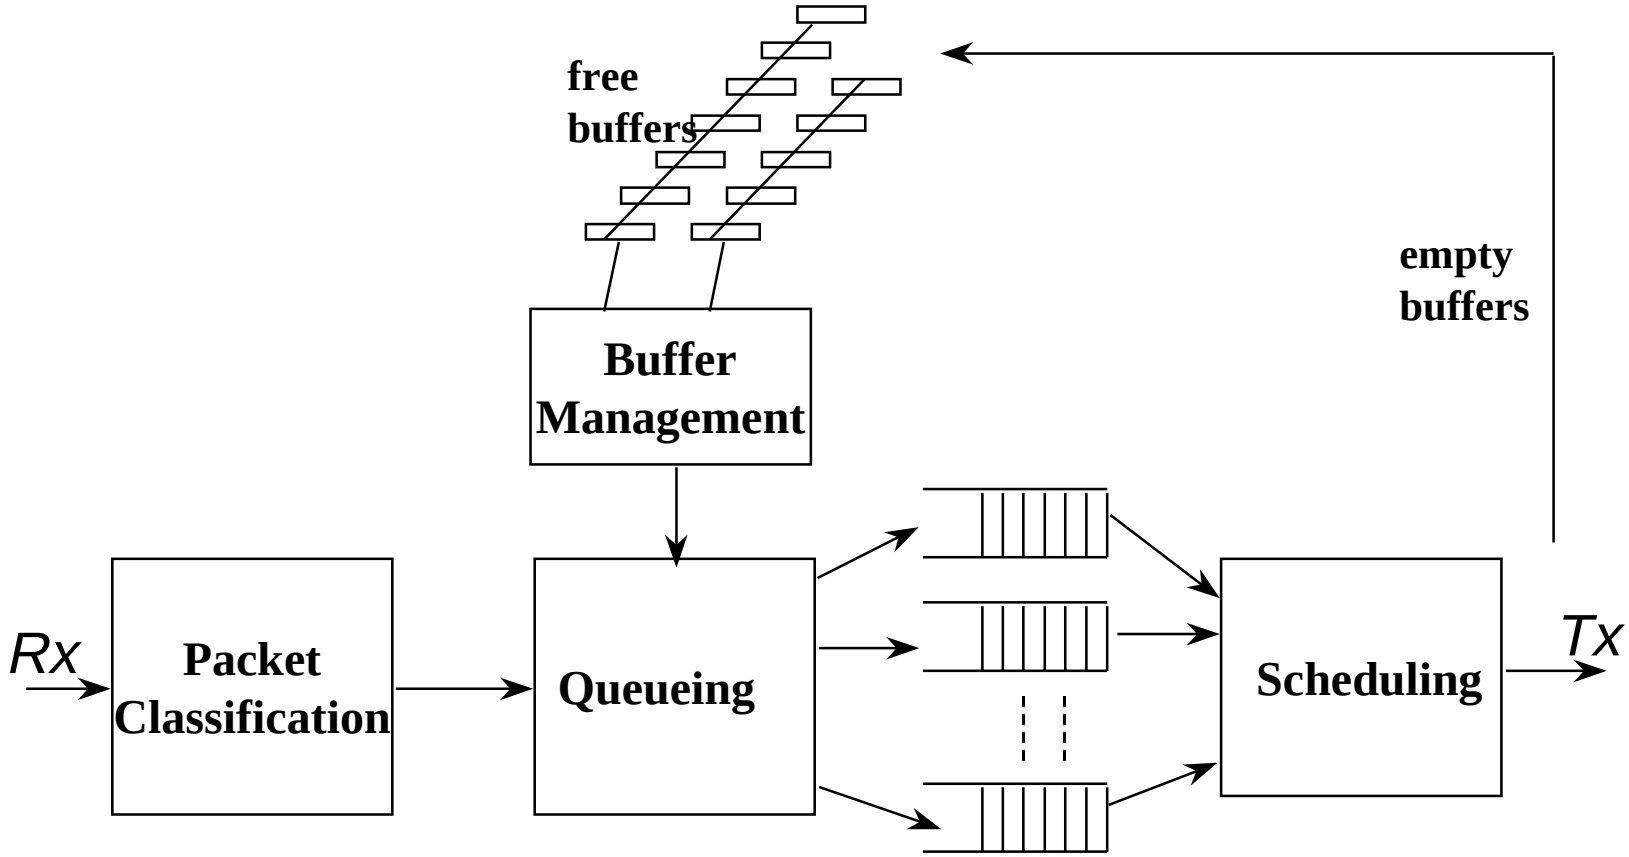

QoS is control of critical resources

- **Buffer management provides control over memory resources**
  - Permits minimum buffer allocation and maximum buffer limits per queue
- **Queuing structure and transmit scheduler provide control over bandwidth resources**
  - Priority queueing assures that high priority traffic gets through the network.
  - Fair queueing assures that congestion in one department or subnet does not interfere with others.
  - Bandwidth controls enable guaranteed service levels.

Conclusion

In the last year products have been developed and shipped that dissolve the myths regarding

Gigabit Speed Routing

and

Quality of Service on Ethernet

extreme networks