

CMSC 714
Lecture 9
Ethernet and Infiniband

Alan Sussman

Notes

- OpenMP assignment due week from Monday
 - don't compile for profiling (with `-pg`) and for OpenMP at same time
 - other questions?
- Reading for next week posted
 - Don't forget to send questions

Ethernet

- One of the first, and definitely the most successful, local area network (LAN) protocol and implementation
- Notice that also targeted at multiprocessing, not just distributed computing
- CSMA/CD wire protocol – carrier sense multiple access, with collision detection
 - a shared physical medium, even with repeaters
 - randomized exponential backoff after a collision
 - error detection using checksums on medium, still need end-to-end error detection since whole packets can be lost
- Modern Ethernet protocols are different – for a switched physical medium, to scale better
 - original at 3Mb/s (shared), 1Gb/sec (switched) common now, with 10Gb/sec and 100Gb/sec in HPC environments

Infiniband

- Designed to support I/O and network connectivity, from a single PCB to a cluster network to a LAN
 - over copper (twisted pairs) and fiber
- Targeted at cluster networks, SANs, and even embedded systems
 - scalable, and provides RAS – “bandwidth out of the box”
 - idea is to extend the on-processor I/O bus to off-chip network
- Switched point-to-point I/O fabric
 - endpoints (host machines, I/O devices, ...) connect to switches, which route connections to other endpoints
 - link speed from 2.5Gb/sec (1X) to 30Gb/sec (12X) by adding more wires – parallel transfers – newer standards use higher link speeds for higher transfer rates
- Protocols described in terms of standard network layers
 - physical, link, network, transport

Infiniband Layers

- Physical

- defines electrical and mechanical characteristics – cables, connectors, pins, etc.

- Link

- packet layout - management and data
- switching - uses local IDs in Local Route Header of a packet
- QoS through Virtual Lanes
- credit based flow control
- data integrity – error correction both for each link (VCRC) and end-to-end (ICRC)

- Network

- route packets across subnets – uses IPv6 addresses (128 bits) in Global Route Header of a packet

Infiniband Layers (cont.)

- **Transport**

- in-order packet delivery – sequence numbers
- segmenting data into packets
- channel multiplexing
- transport services – reliable/unreliable connection/datagram
- all implemented in hardware