High-performance IP Forwarding Using Host Interface Peering

Joe Touch, touch@isi.edu
Anne Hutton, hutton@isi.edu
Simon Walton, simonw@isi.edu
Stephen Suryaputra, surya@isi.edu
USC/ISI Computer Networks Division

This work is supported by the Defense Advanced Research Projects Agency through Ft. Huachuca contract #DABT63-93-C0062 entitled "Netstation Architecture and Advanced Atomic Network".The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Department of the Army, the Defense Advanced Research Projects Agency, or the U.S. Government.

Host-Based Forwarding

Benefits

- Programmability
- Commodity Platforms and Network Interfaces
- Network Interface Cards track technology advances
- NICs precede line cards (if line cards exist at all).

Example uses in Research

- in testbeds DARTnet and successor CAIRN
- Active Networks.
- ATOMIC project supporting Myrinet LAN at ISI

Problems

Latency

- Store and forward copying

Bandwidth

- bus limited backplane

CPU load

- interrupts
- cycles to manage transfers (PIO more than DMA)

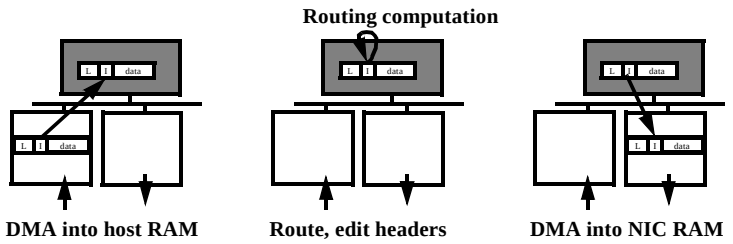

Conventional Forwarding, with DMA to/from Host Memory

Solution: Forwarding using Peer DMA

Two approaches:

- packet on NIC
- data on NIC, copy header to host

Results:

- UDP thoughput up by over 40%
- CPU pegged for small packet sizes and multiple sources
- Relieves CPU load by 35% for 2 sources
- Max packet per sec. 12,000 @ 128 byte packet sizes
- Worse for TO-Host traffic (PIO)

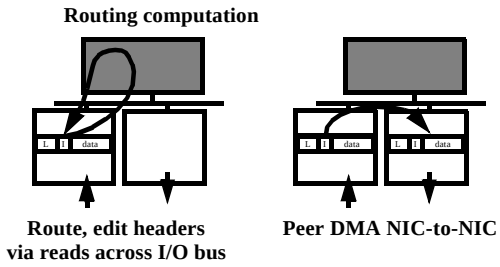

Data and header remain on NIC, followed by peer DMA

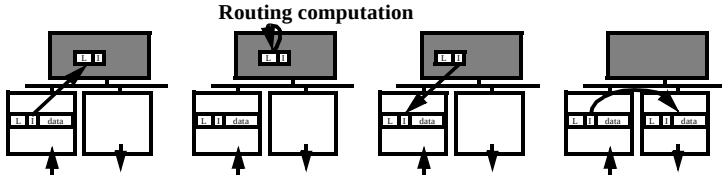

Header Copy, data remains on NIC, followed by peer DMA

Implications for NIC

NIC design

- support DMA
- sufficient shared memory for packet storage
- Co-processor available on NIC ?

Packet issues

- Fragmentation not required or trivial
- Packet data not utilised by CPU (not so in Active Nets)

Host issues

- I/O subsystem supporting DMA

Future Work?

Processing

- How much involvement is needed by Host CPU?
- How much can be done on NIC?

Buffers

- Is buffering required for send and receive on NICs?

Integration

- How to integrate Peer DMA forwarding with TO-Host data?
- Implications for early demux NIC architectures? (APIC)

Throughput Comparison for Three driver Types (UDP)

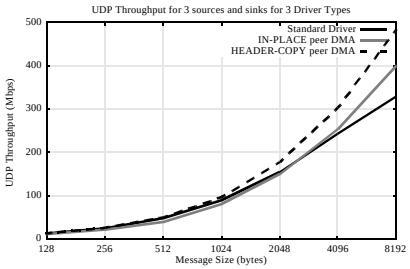

Packet Rate Throughput for Header Copy

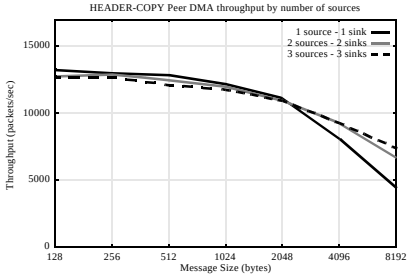

Improvement of peer DMA over Standard Driver