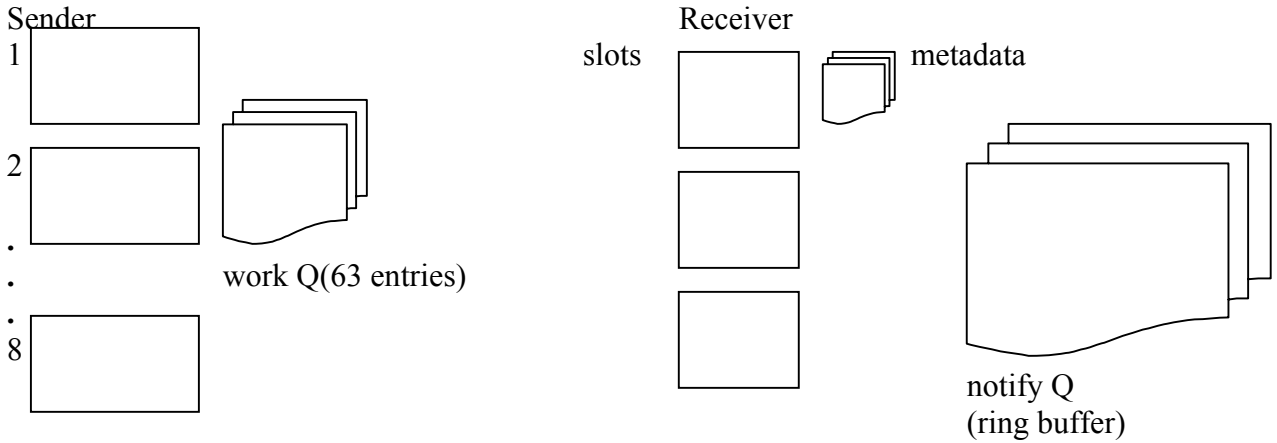


Wednesday, March 13 2002

Note for reading papers:

1. draw the picture that they should have of what they did.
2. 2. Work from there

Hamlyn




slots: pinned areas of virtual memory
metadata: 128B/msg

Usage:

1. Short normal msg

Work Q entry




Header = slot # + offset
Metadata index
Data goes to slot @ offset.


Entry is added to notify Q

2. Long normal msg

Message is composed and sent.



3. a one word msg



Claim:

Hamlyn can do scatter gather. See case 2.

Gather: from send area

Scatter: multiple messages to receive areas, get help from notify mechanism, but only one slot.


Zero-copy:

Remote write lets you write to "right" place if your data is always where it should be.

Application layer is responsible for ACKs READ: buffer management.

Paper doesn't actually say how to USE this.

Application Q




send tail pointer back: windowing protocol.

More probably, a remote DMA read

Thekkath NFS:

Really Looks like client Q → server



Rewrite/write

Send areas contiguous in physical mem?

Yes: then this is all easy.

No: Either DMA has translation & knows about page bounds.

OR HP phys. addr. has VA on IO bus.

Note: PCI does not do this!

