

Performance and ARQ

From before, we know that the utilization of an ARQ protocol is: (C is the network latency, E is the expected number of errors per message and T is the timeout period).

$$U = \frac{M/R}{L/R + C + E(T + L/R)}$$

To simplify, let $M = L$ and divide through as before:


$$U = \frac{1}{1 + C/(L/R) + E(T/(L/R) + 1)}$$

If the probability of a successful send is high, which we expect, then the latter term drops out:

$$U = \frac{1}{1 + C/(L/R)}$$

- Efficient use of the channel is inversely proportional to the ratio of the network latency to the transmit time.
- High utilization translates into high throughput and low response time.
- Normally, network latency is much larger than transmit time.
- There is little that can be done to modify C.

Pipelining Messages



You might send N messages instead of one while waiting for an ACK:

$$U = \frac{N \cdot \text{Data Time}}{N \cdot \text{Message Time} + \text{Network Latency}}$$

Using the equation above with L/R replaced by NL/R ,

$$U = \frac{1}{1 + C/(NL/R)},$$

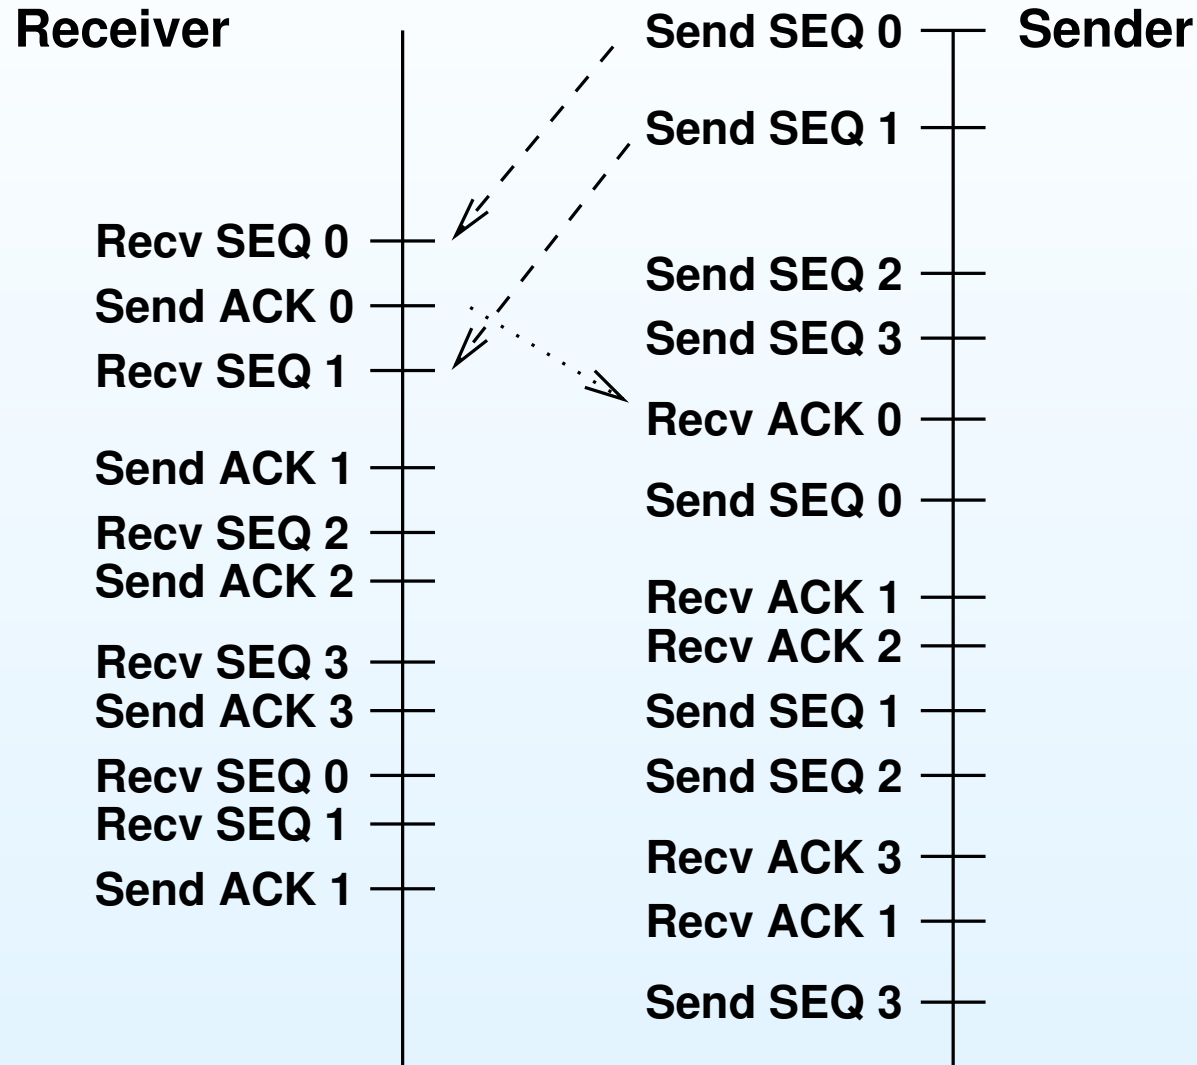
we see that utilization gets higher as N increases. But there is no increase beyond $NL/R = C$. You pay the price of C for the first packet, but after that, you would theoretically have an ACK immediately after each message was sent, so the utilization approaches 1.0 as N gets large. But only N messages can be in the channel at any given time if $NL/R = C$.

Sliding Window Protocols

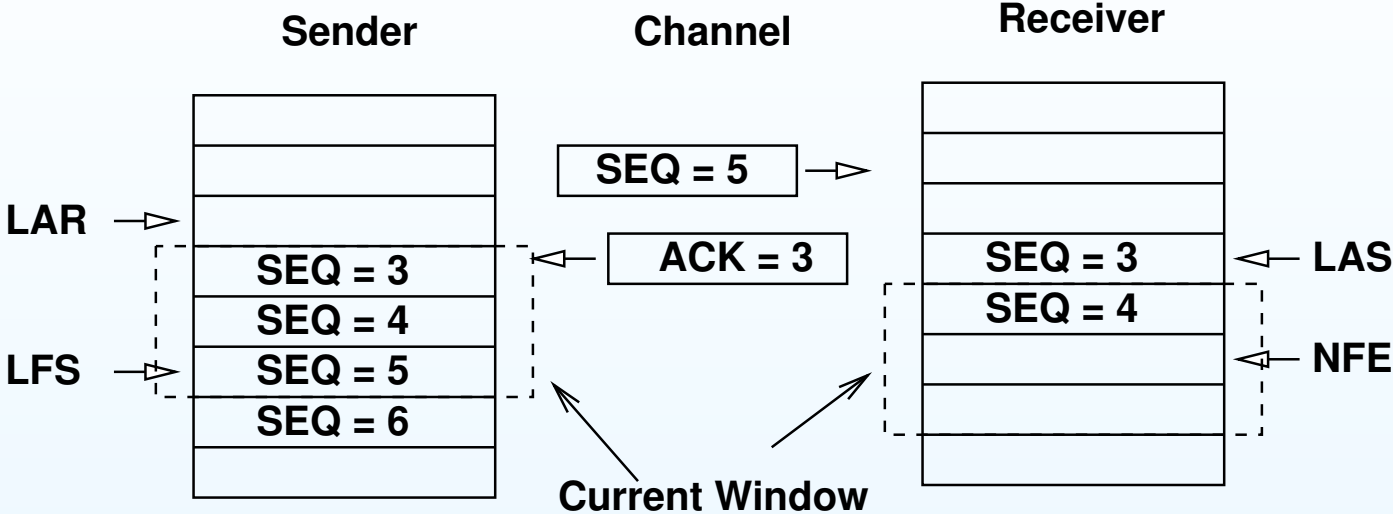
Sliding Window Protocols (SWP's) are ARQ protocols that use pipelining:

- The window defines the number of unacknowledged messages that can be in the channel at any one time.
- Each message must be acknowledged, although not independently.
- The limiting factor might be buffer space if network latency is very high.
- SWP's are connection-oriented to allow protocol parameters to be negotiated.
- Most SWP's are reliable, except for delayed packets.
- SWP's are usually duplex.

2-bit Sequence Number Example



Practical SWP's

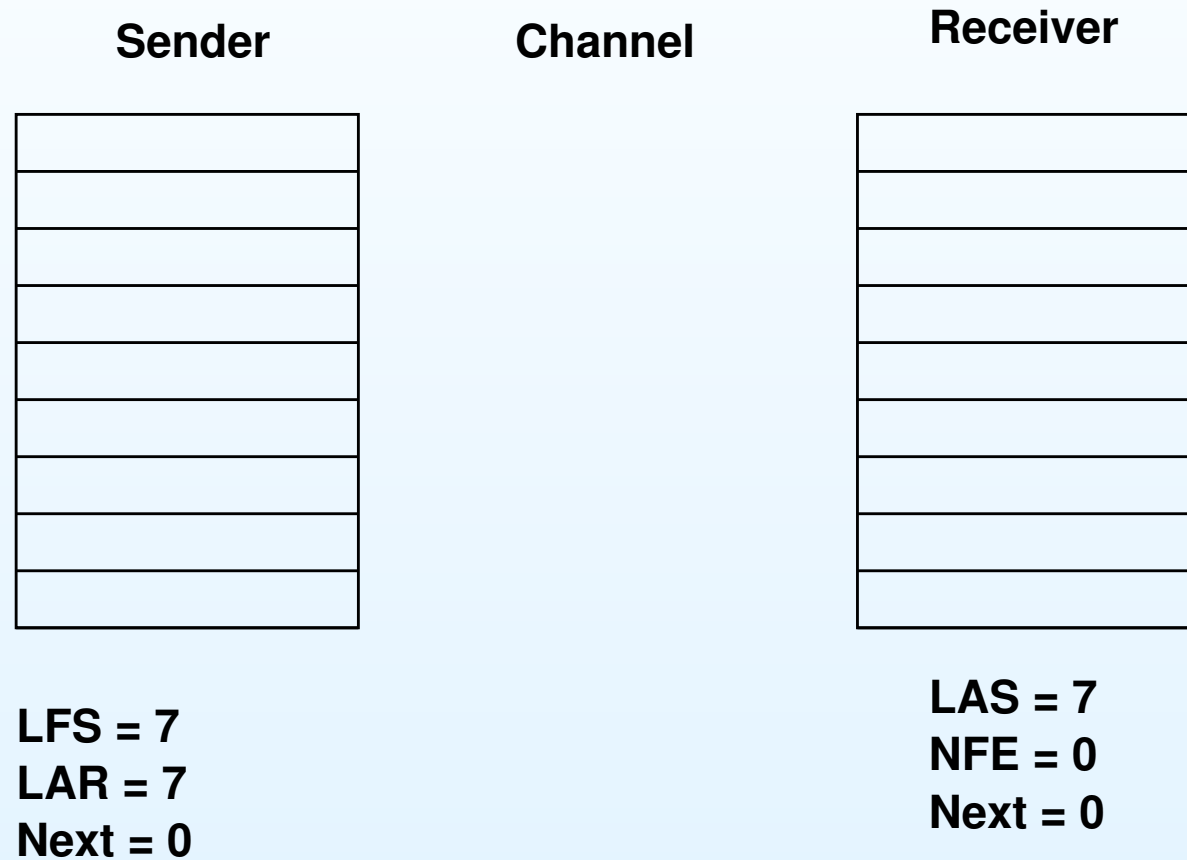


**A 3-bit Sliding Window Protocol
with Window Size = 3**

- Receivers ACK with NFE.
- An ACK acknowledges all sequence numbers up to, but not including, that sequence number.
- Receivers send an ACK for any message they receive, even if it is a duplicate, to insure that the sender knows their current state.
- The sender has N buffers which limit the number of outstanding messages, but the sequence numbers can have a larger range.

3-bit Sequence Number Example

An SWP uses 3-bit sequence numbers, has 9 buffers for messages and has a current maximum window size of 5.



An SWP Example

Practicalities

- Flights and acknowledgements
- The impact of a lost packet or ACK.
- Error handling
 - Go-Back-N
 - Selective Repeat
 - Go-Back-N effectively means that messages must arrive in order, as

Byte-oriented SWP's

- Because it is unreasonable to expect applications to use fixed length messages.
- And it provides excellent control of the channel.
- The window is the number of bytes currently in the channel.