




Reliable Data Transfer




Reliable Data Transfer




Reliable Data Transfer




Reliable Data Transfer




Reliable Data Transfer



Reliable Data Transfer




Reliable Data Transfer




Focus only on packet-delivery problems

Reliable Data Transfer




Focus only on packet-delivery problems

Reliable Data Transfer




Focus only on packet-delivery problems

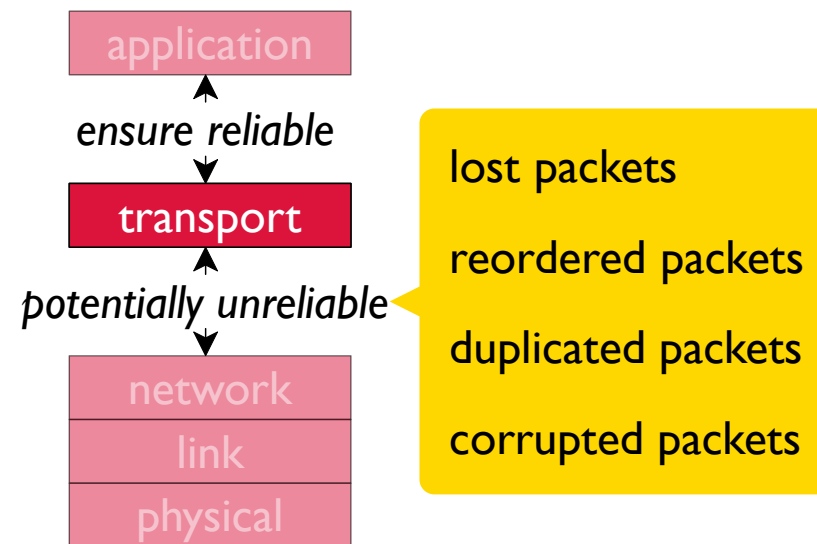
Reliable Data Transfer




Reliable Data Transfer



Reliable Data Transfer




Reliable Data Transfer




Reliable Data Transfer

“reliable data transfer”




Reliable Data Transfer

“reliable data transfer”




“unreliable data transfer”

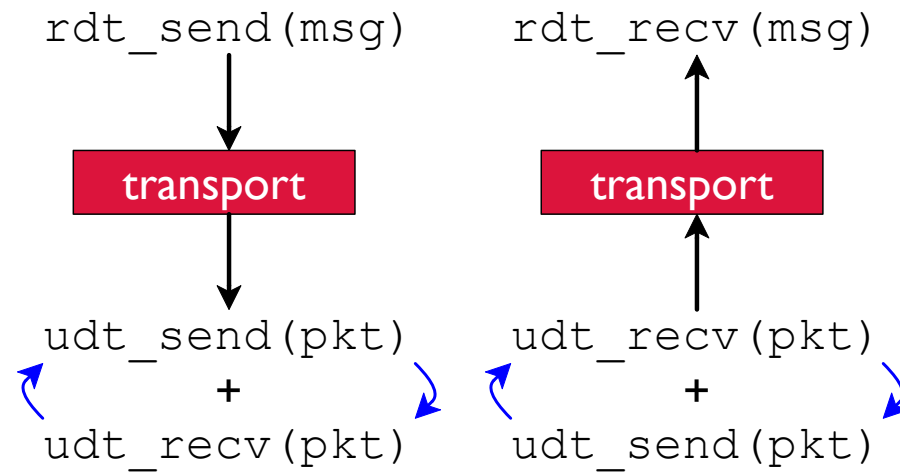
Reliable Data Transfer




Reliable Data Transfer



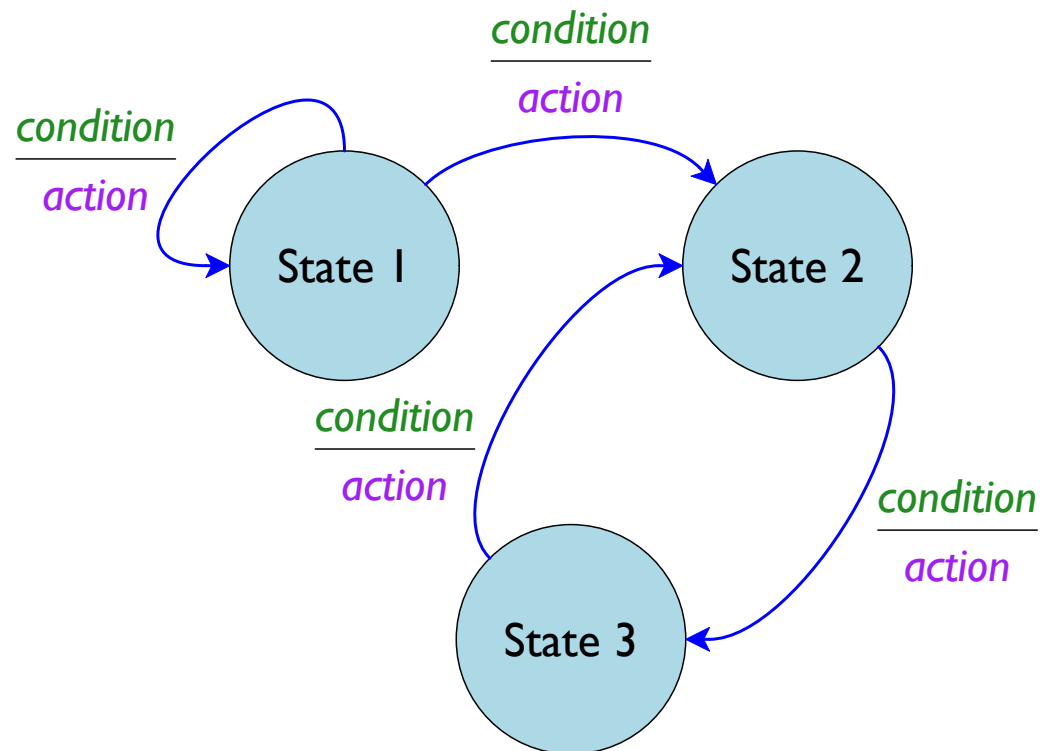
Reliable Data Transfer



State Machines




State Machines




State Machines

Vending machine example



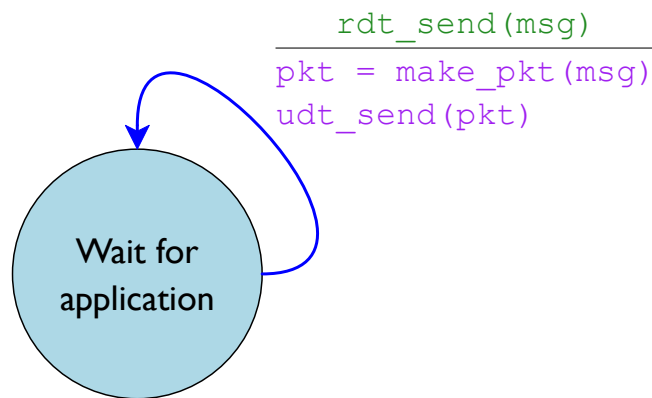
Assuming Reliable `udt_send` and `udt_rcv`

sending host

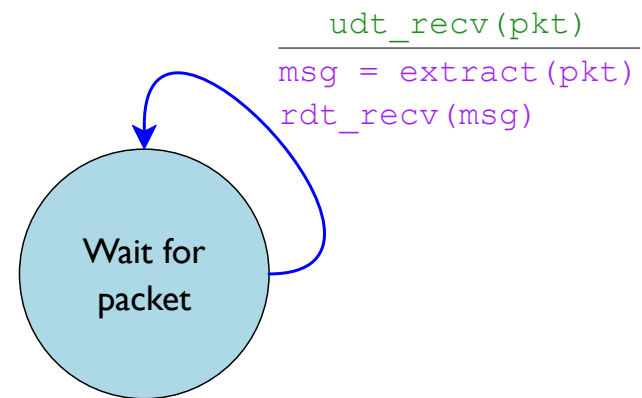


Assuming Reliable `udt_send` and `udt_rcv`

sending host



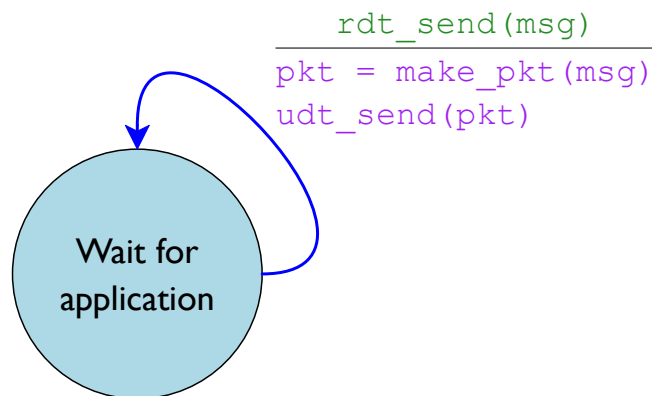
receiving host



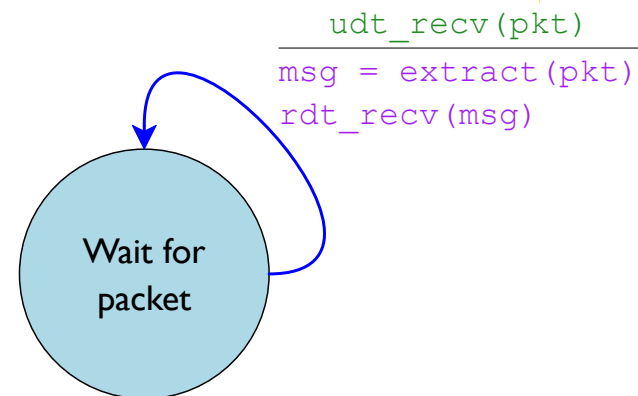
Assuming Reliable `udt_send` and `udt_rcv`

What if `pkt` is corrupted?

sending host



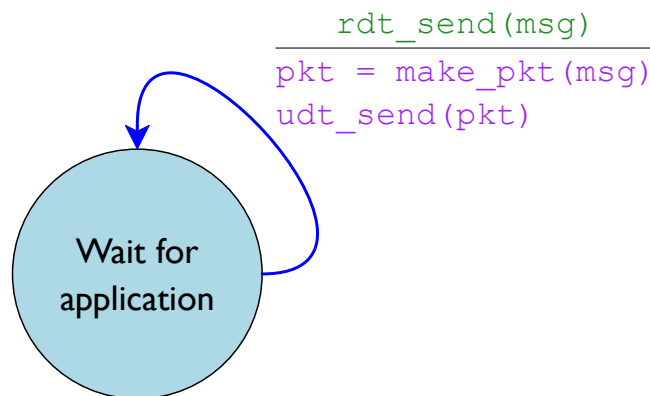
receiving host



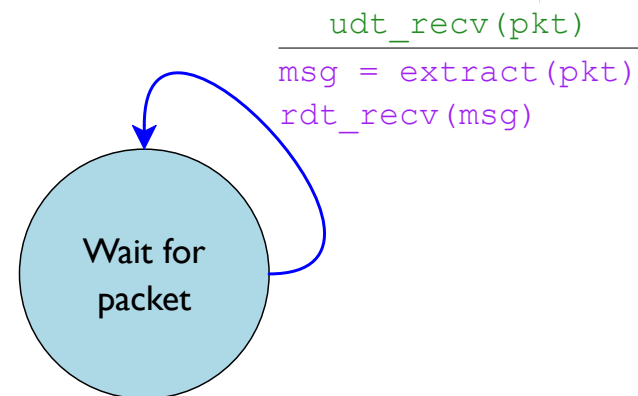
Assuming Reliable `udt_send` and `udt_rcv`

What if `pkt` is corrupted?

sending host




receiving host




Send twice and check as the same?

Checksum



Checksum

I'd like 1 apple, 2 bananas, and 3 cherries




Checksum

I'd like 1 apple, 2 bananas, and 3 cherries




Ok: 1 apple, 2 bananas, and 2 cherries



Checksum


I'd like 1 apple, 2 bananas, and 3 cherries



Ok: 1 apple, 2 bananas, and 2 cherries

Checksum

I'd like 1 apple, 2 bananas, and 3 cherries — which is 6 total



Checksum


I'd like 1 apple, 2 bananas, and 3 cherries — which is 6 total



Ok: 1 apple, 2 bananas, and 2 cherries — which is 6 total

Checksum

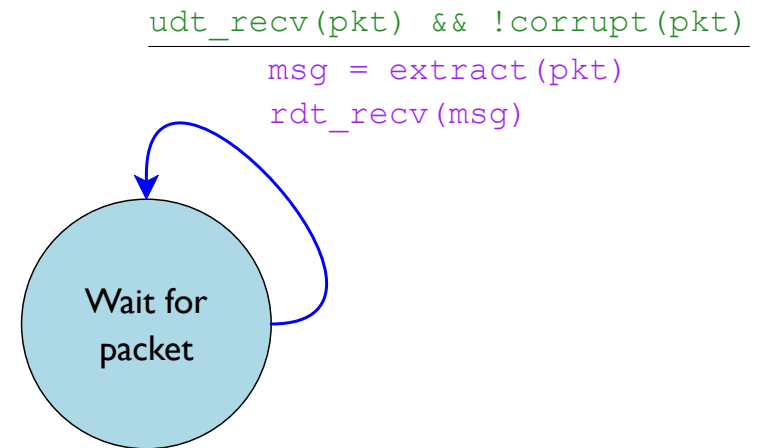

I'd like 1 apple, 2 bananas, and 3 cherries — which is 6 total



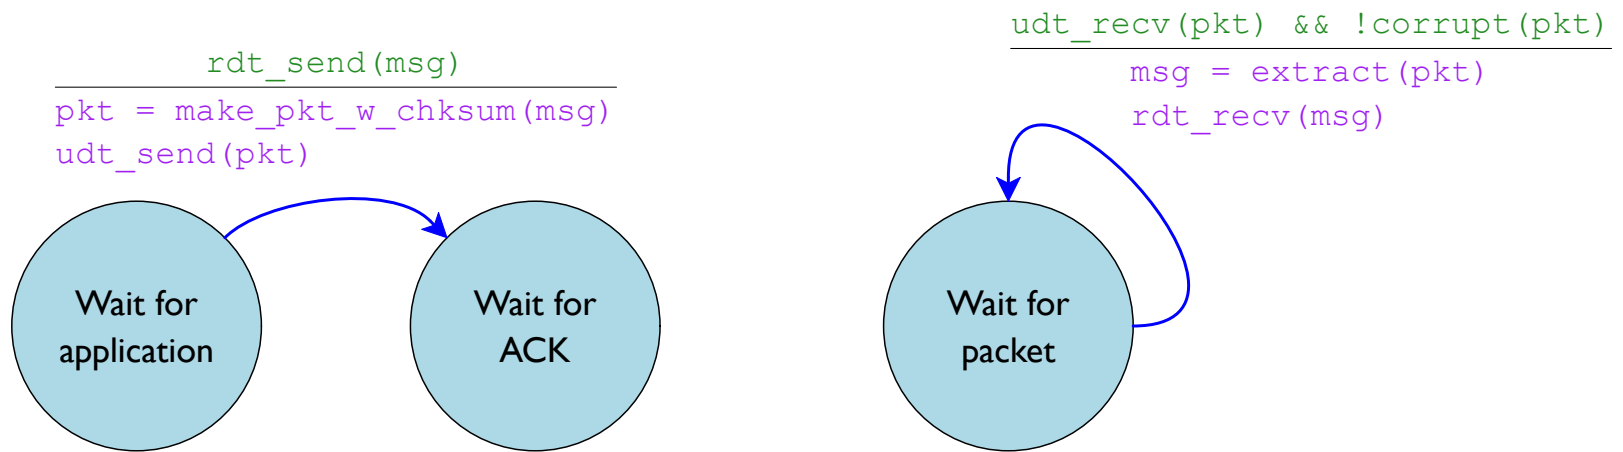
Ok: 1 apple, 2 bananas, and 2 cherries — which is 6 total

To deal with lots of numbers, just keep low bits of the sum

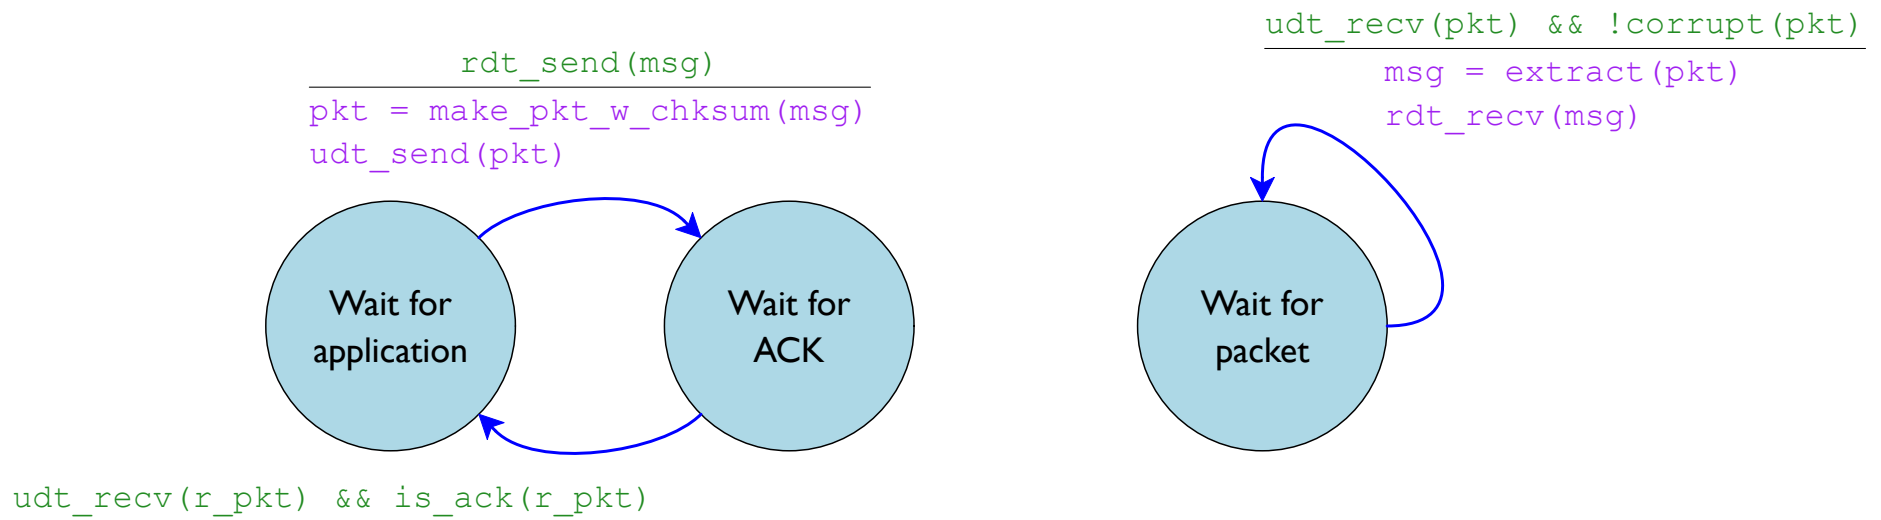
Using a Checksum



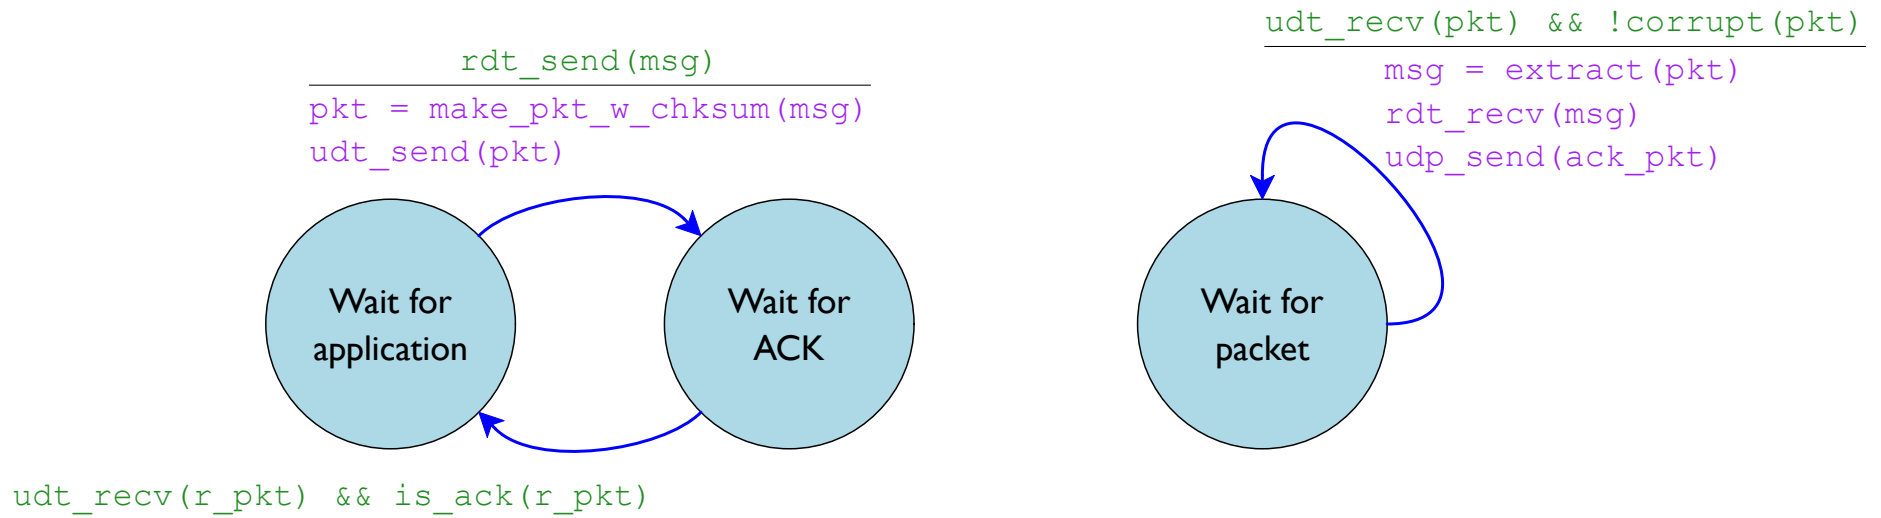
Using a Checksum



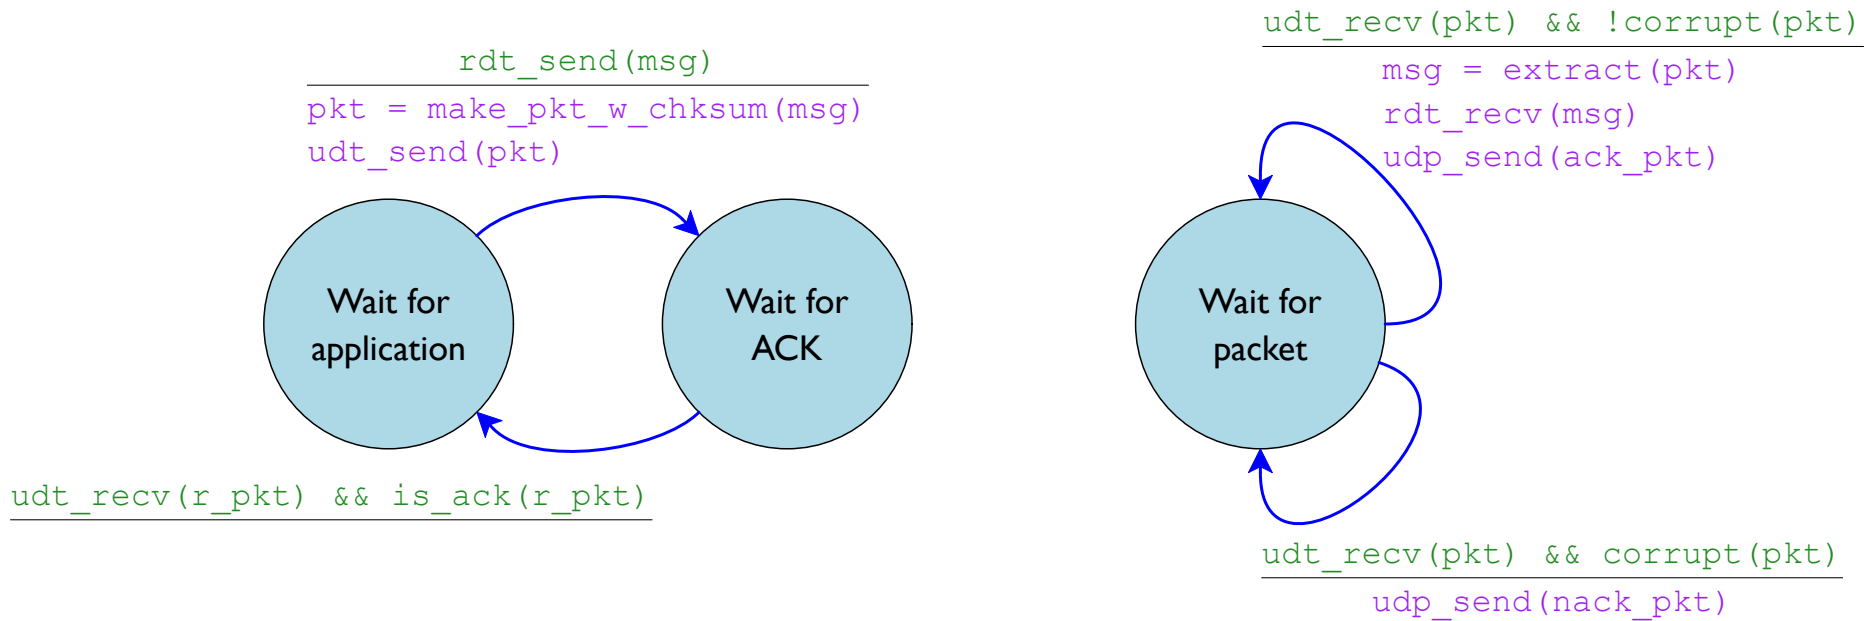
Using a Checksum



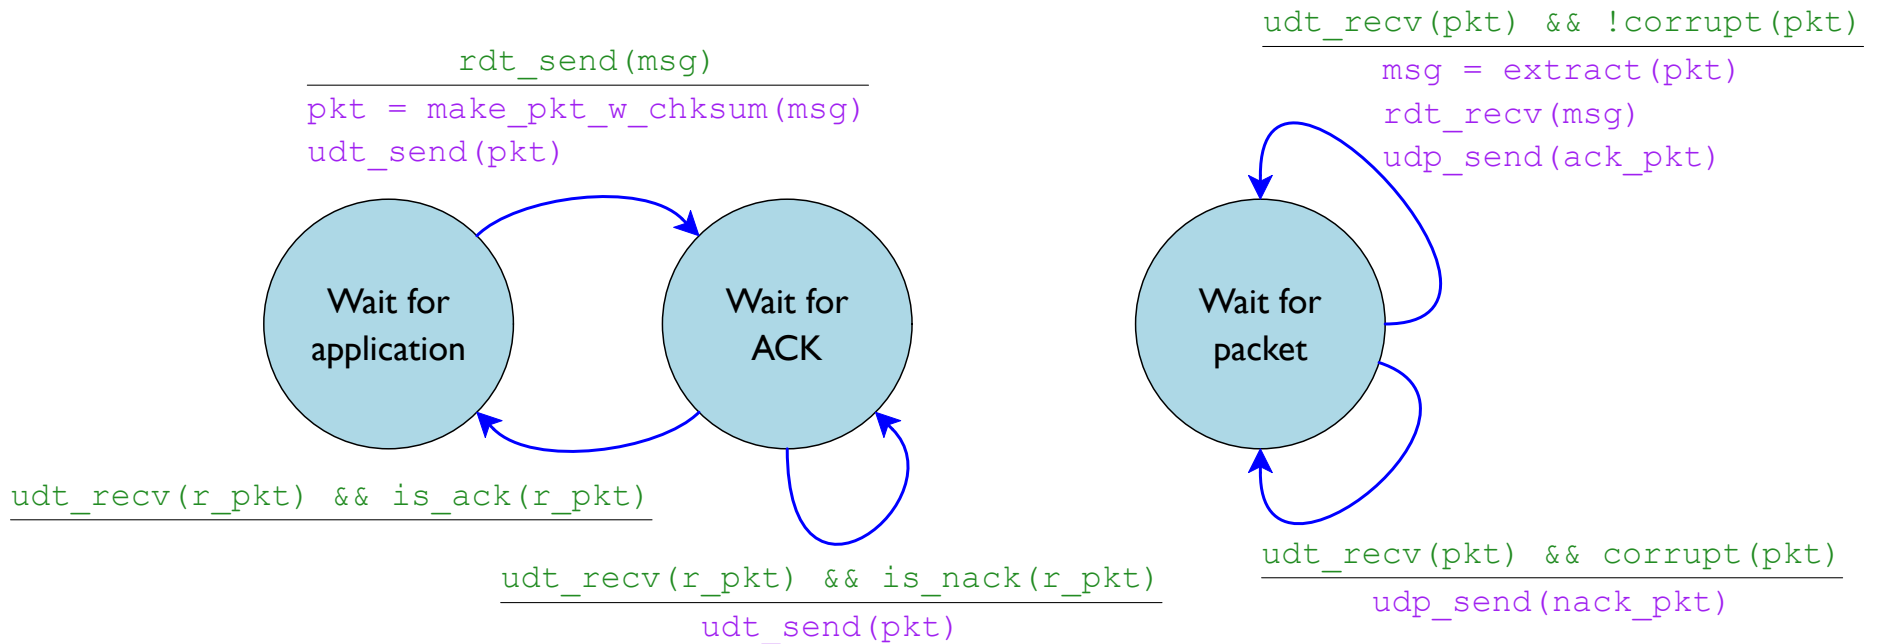
Using a Checksum



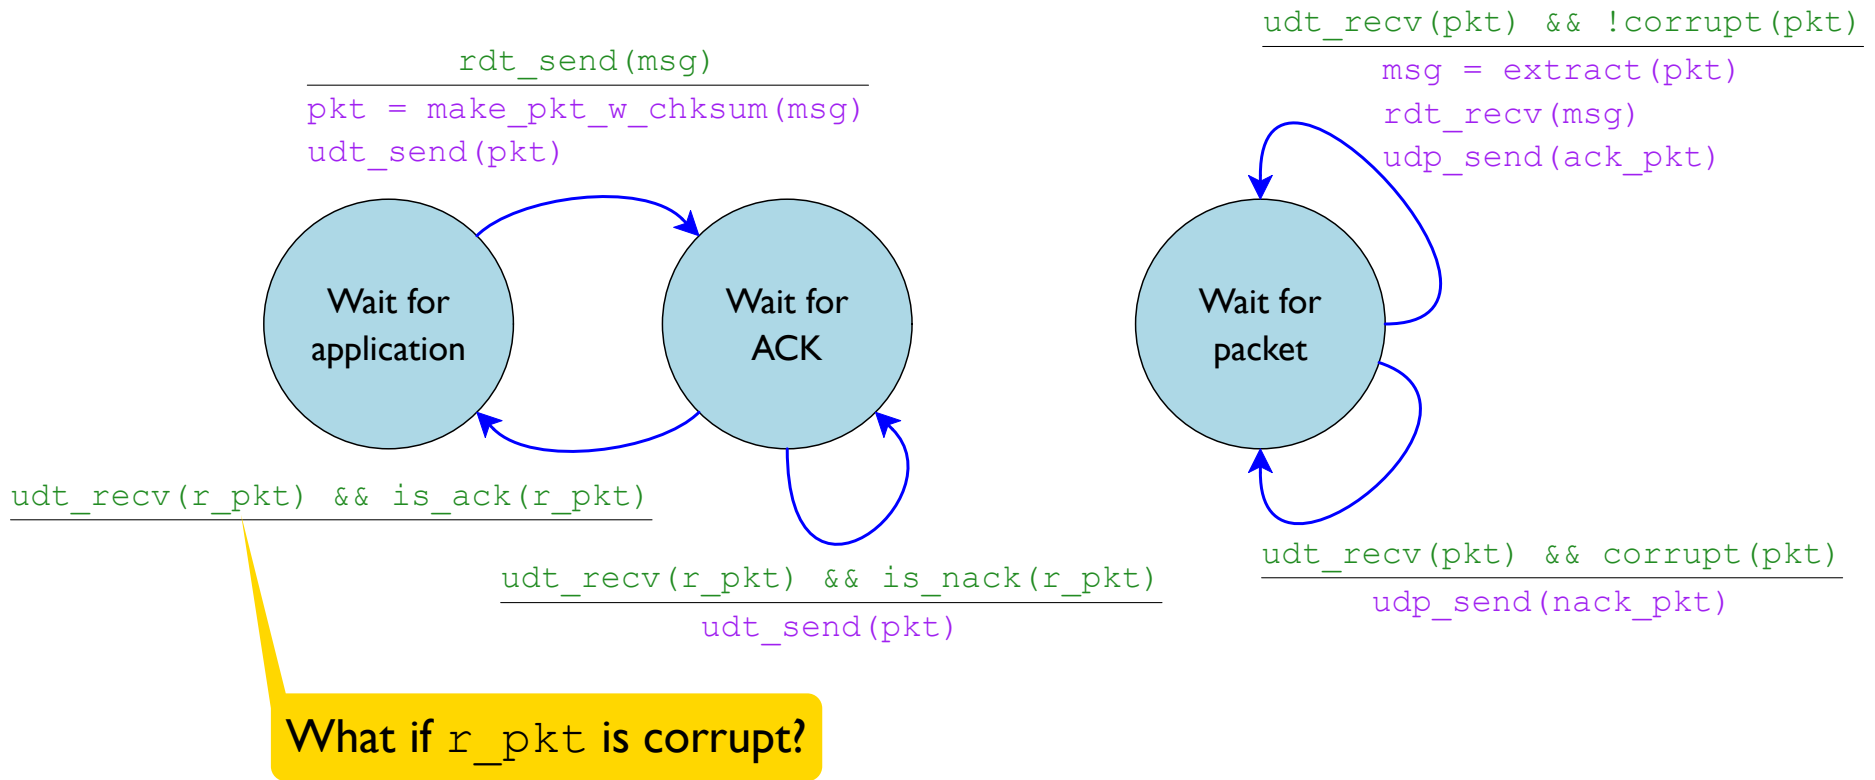
Using a Checksum



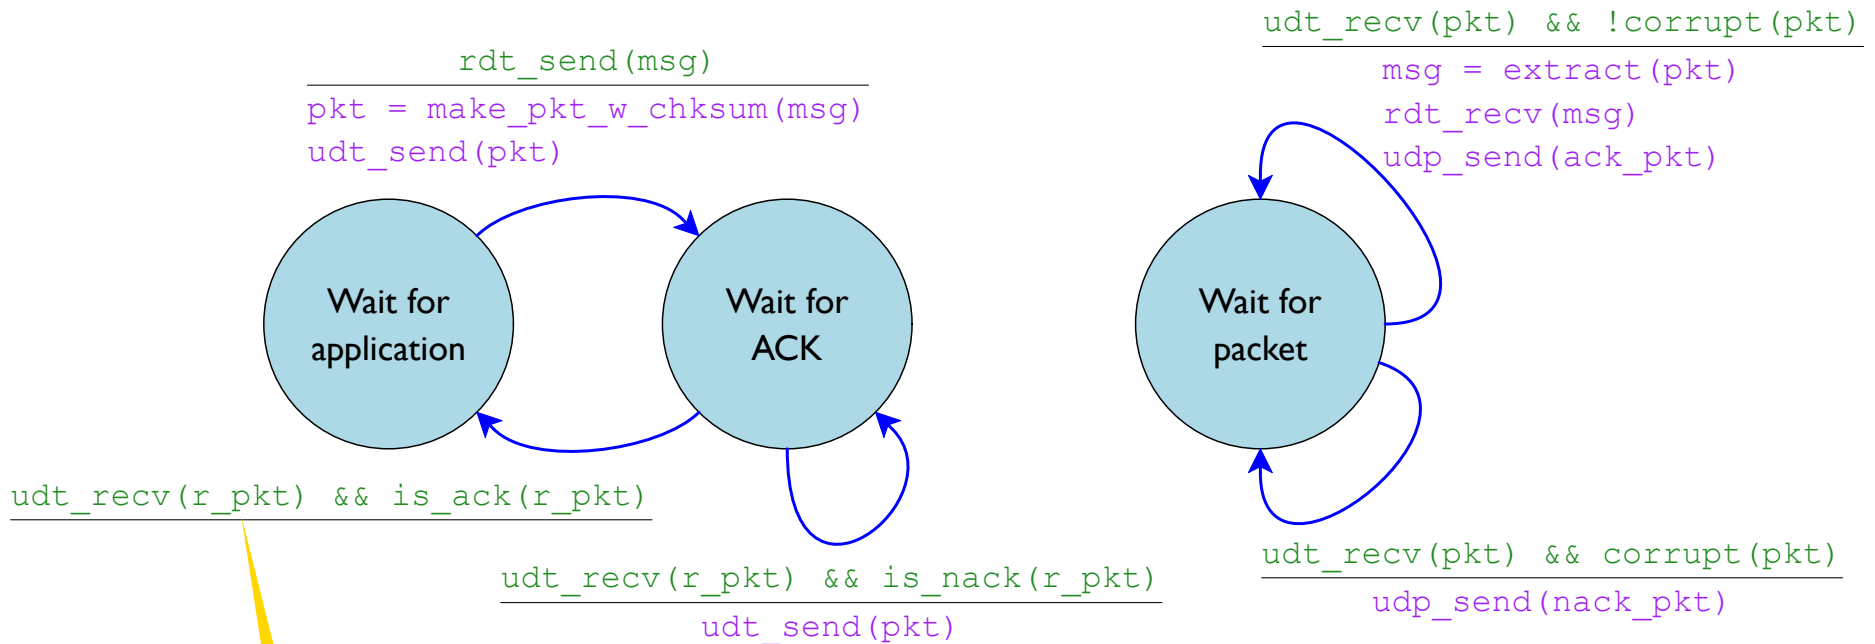
Using a Checksum



Using a Checksum

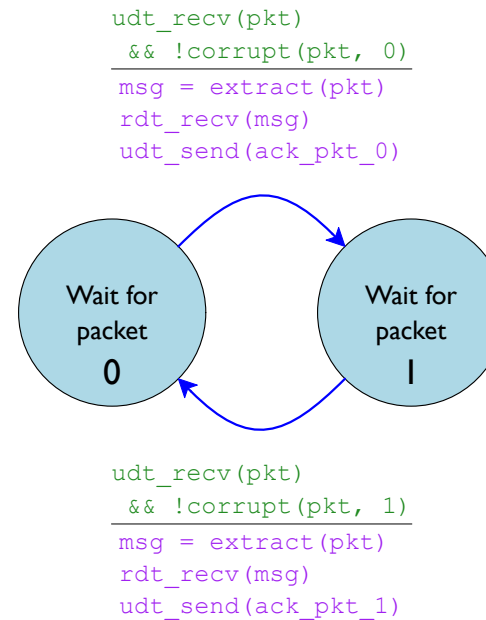
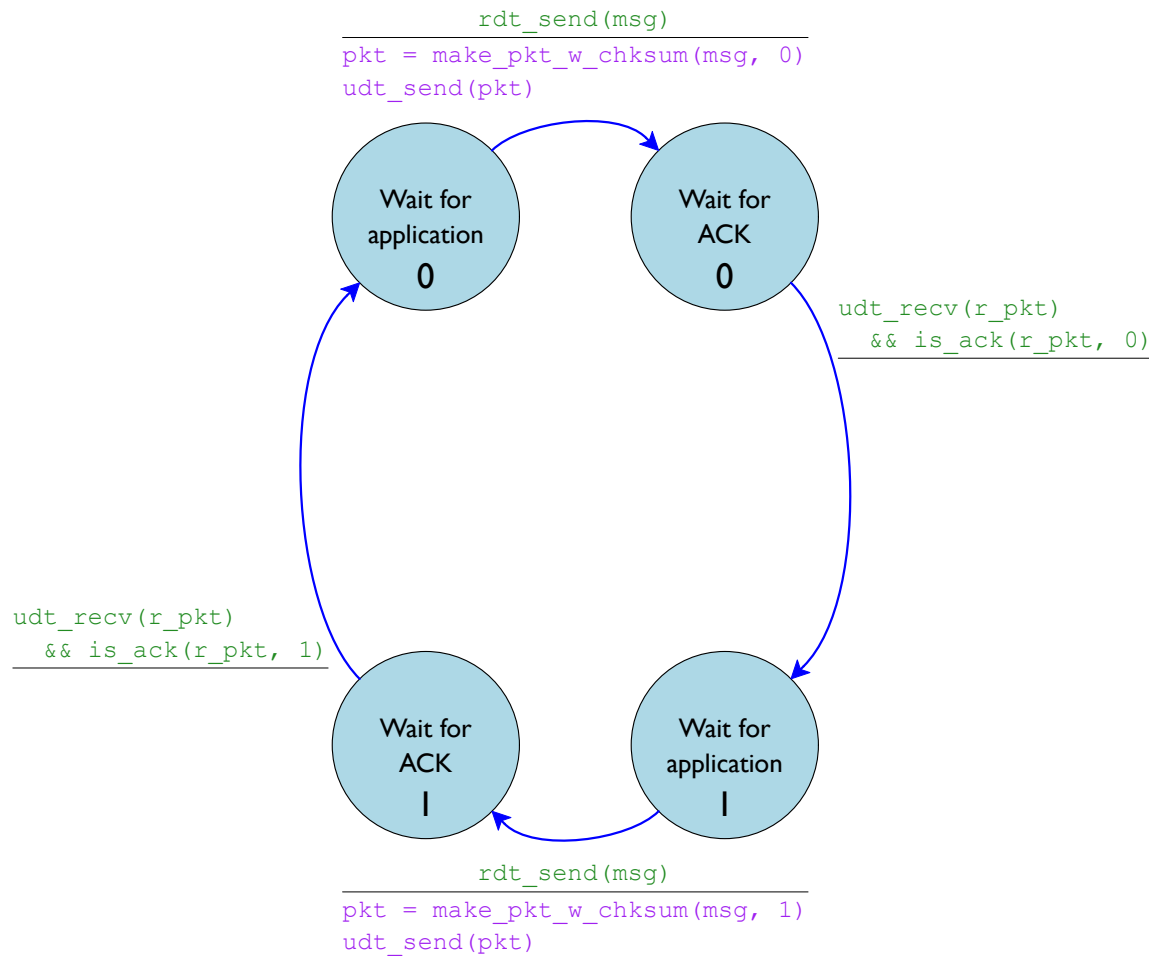


Using a Checksum

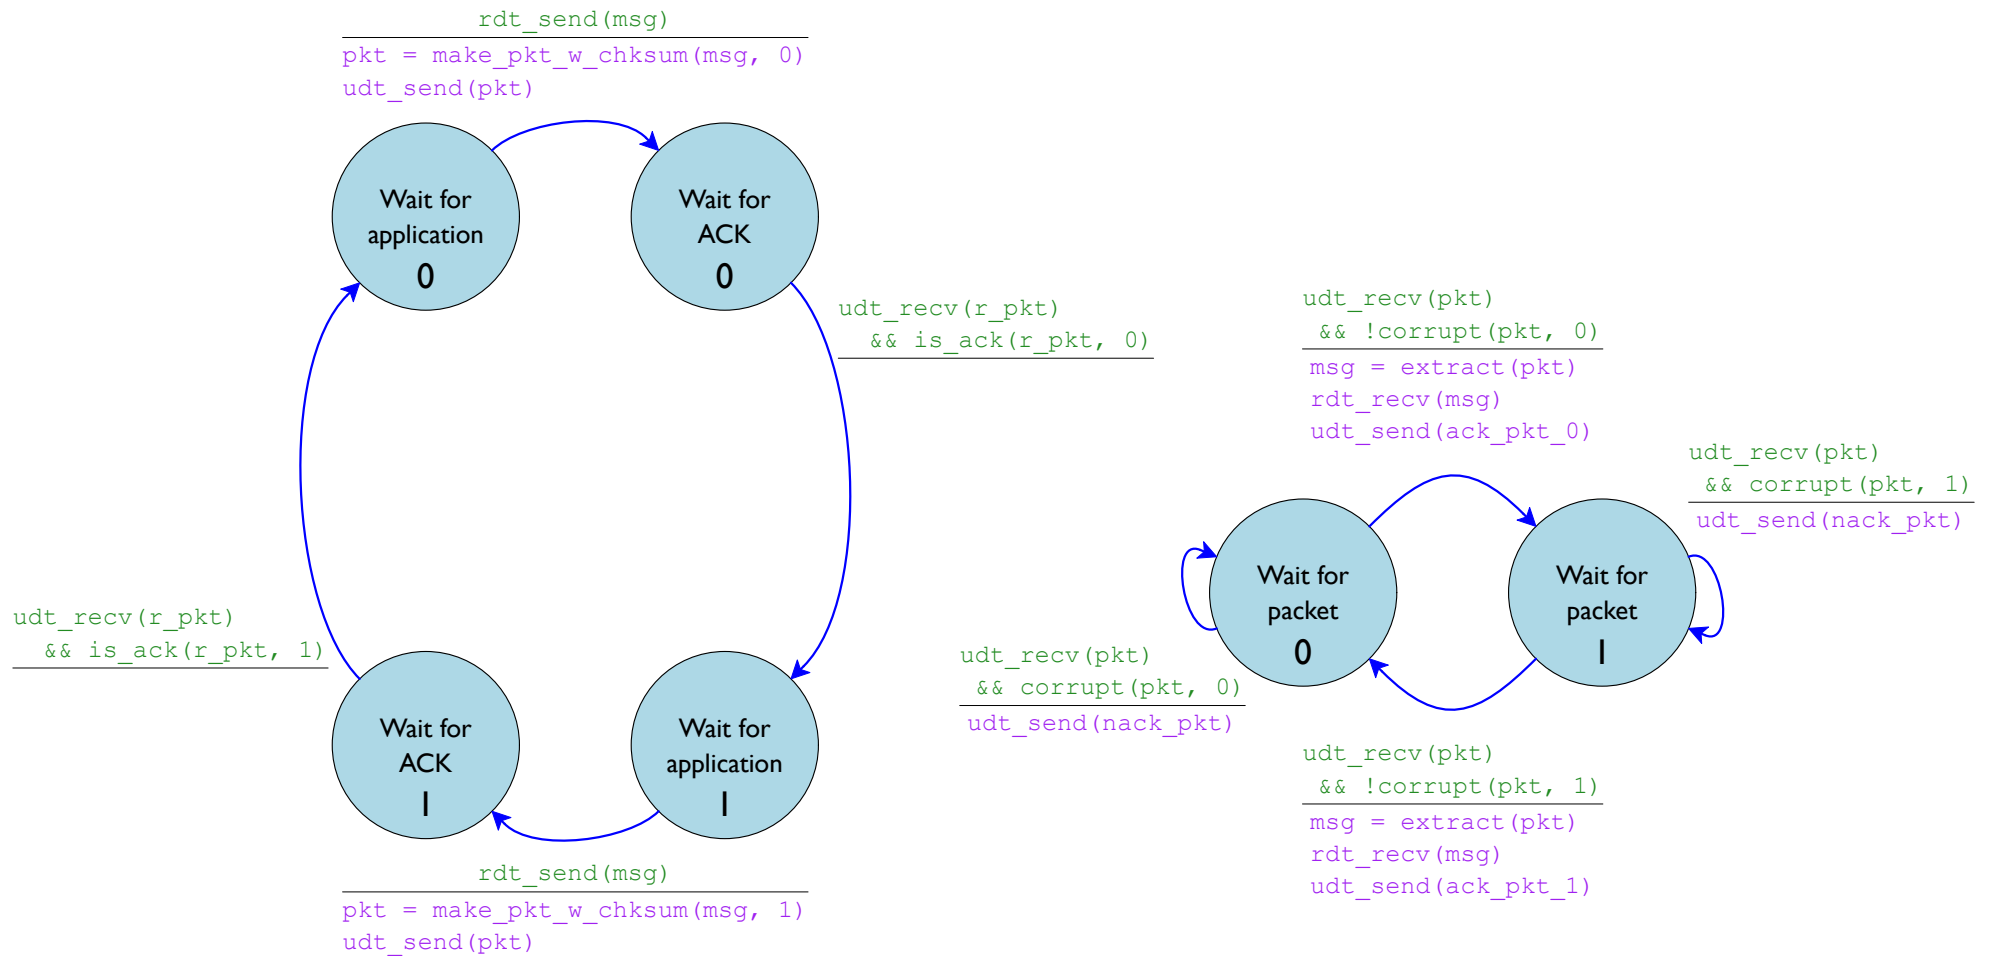


How do we avoid an ACK of ACK of ACK...?

Handling ACK Corruption

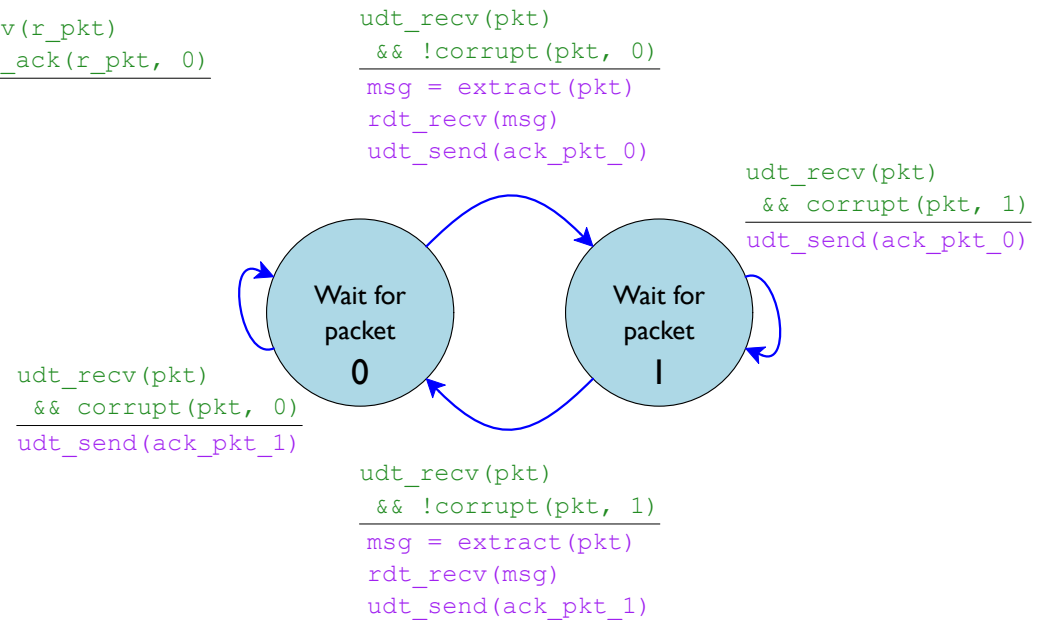
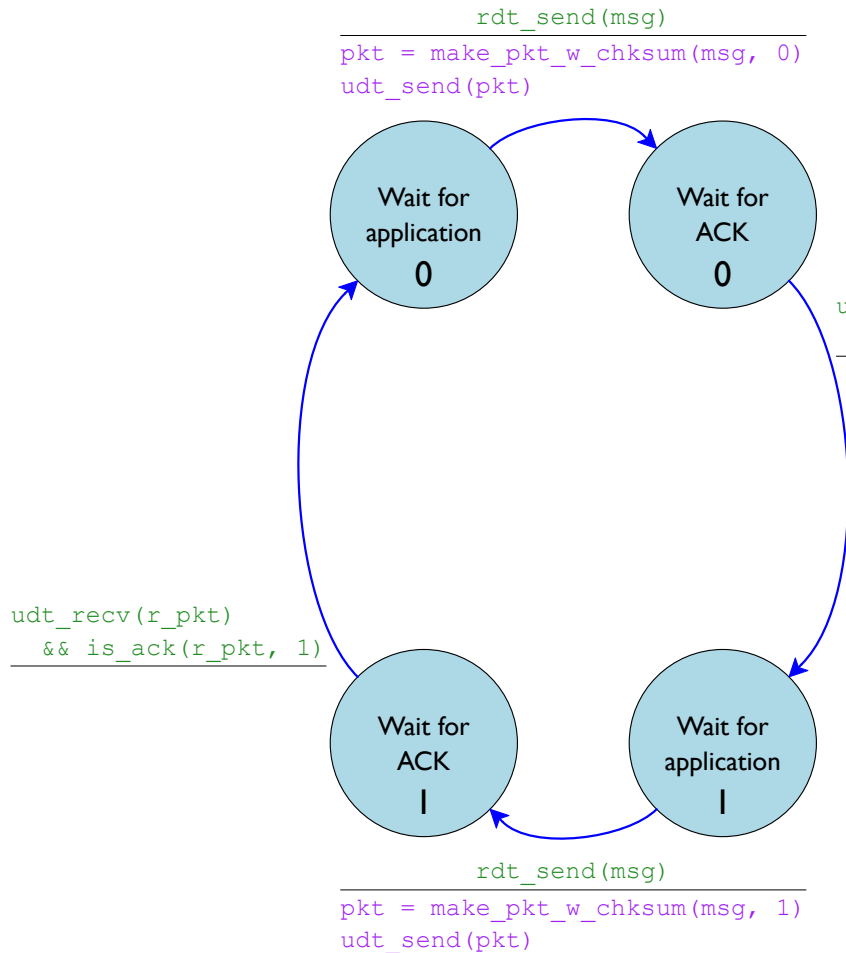


Handling ACK Corruption

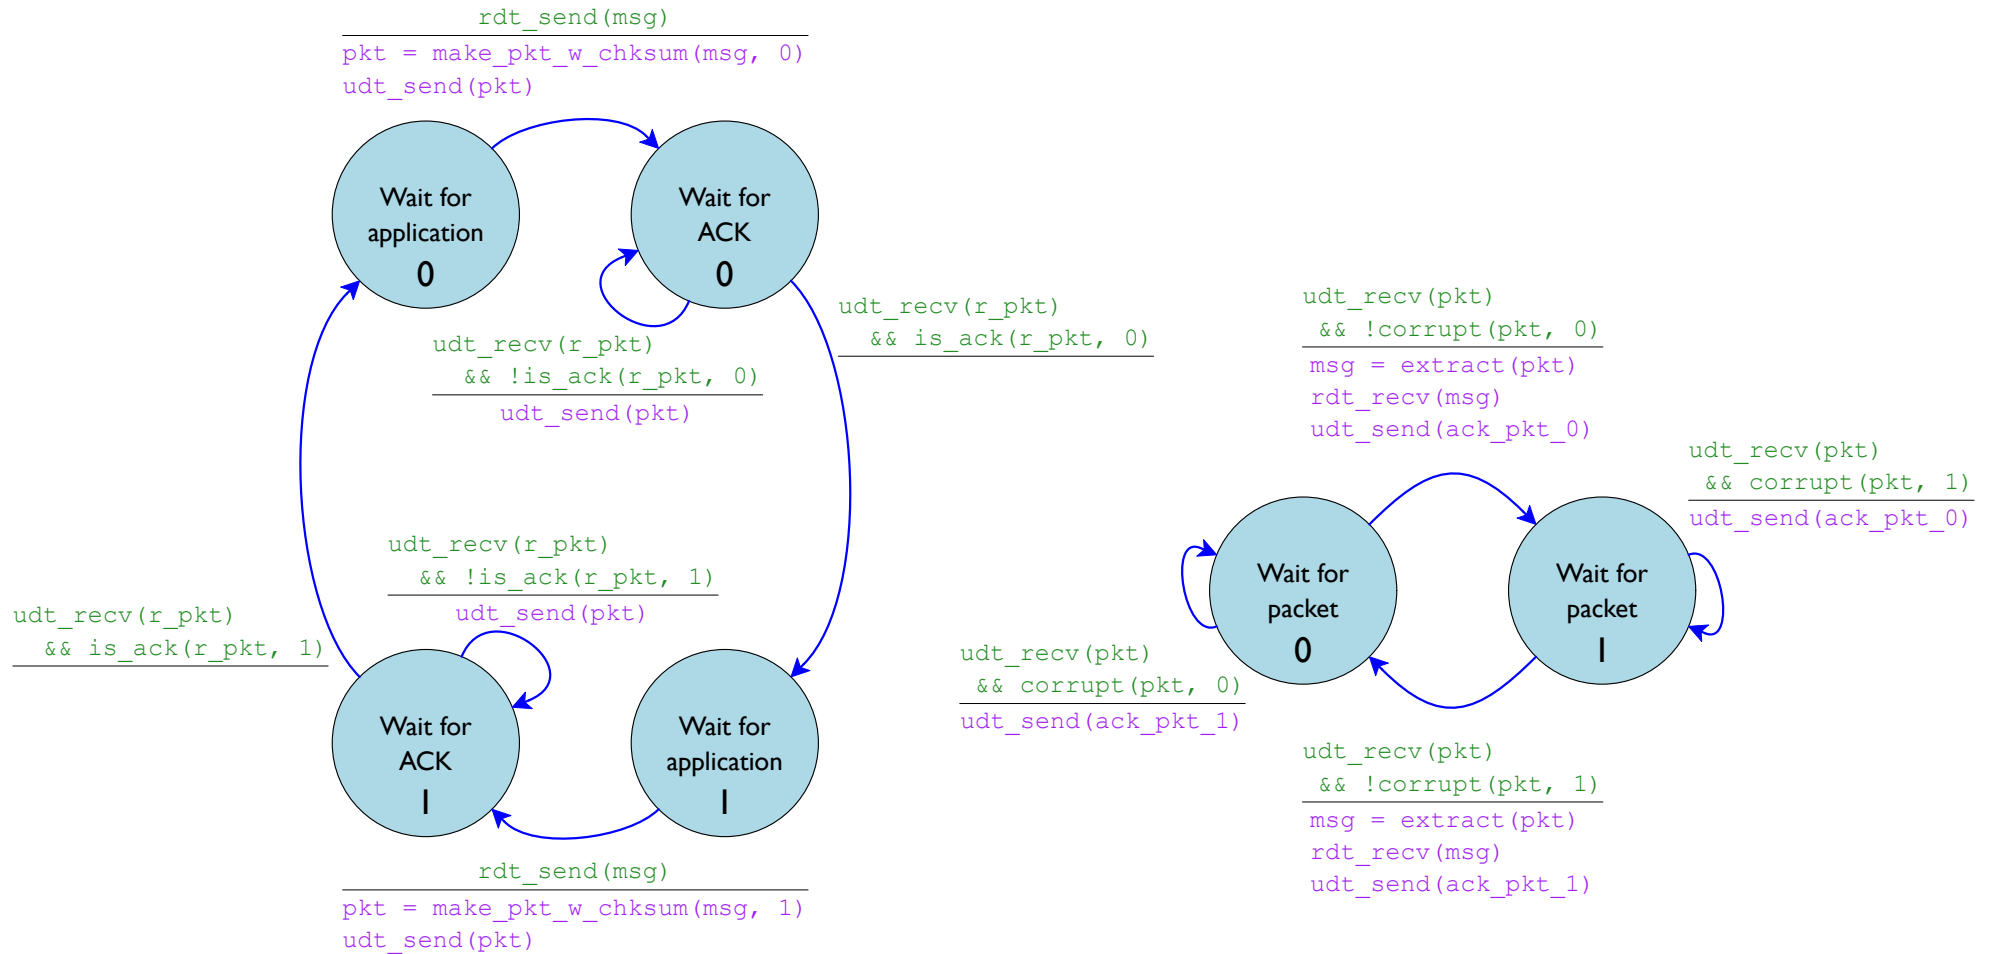


Handling ACK Corruption

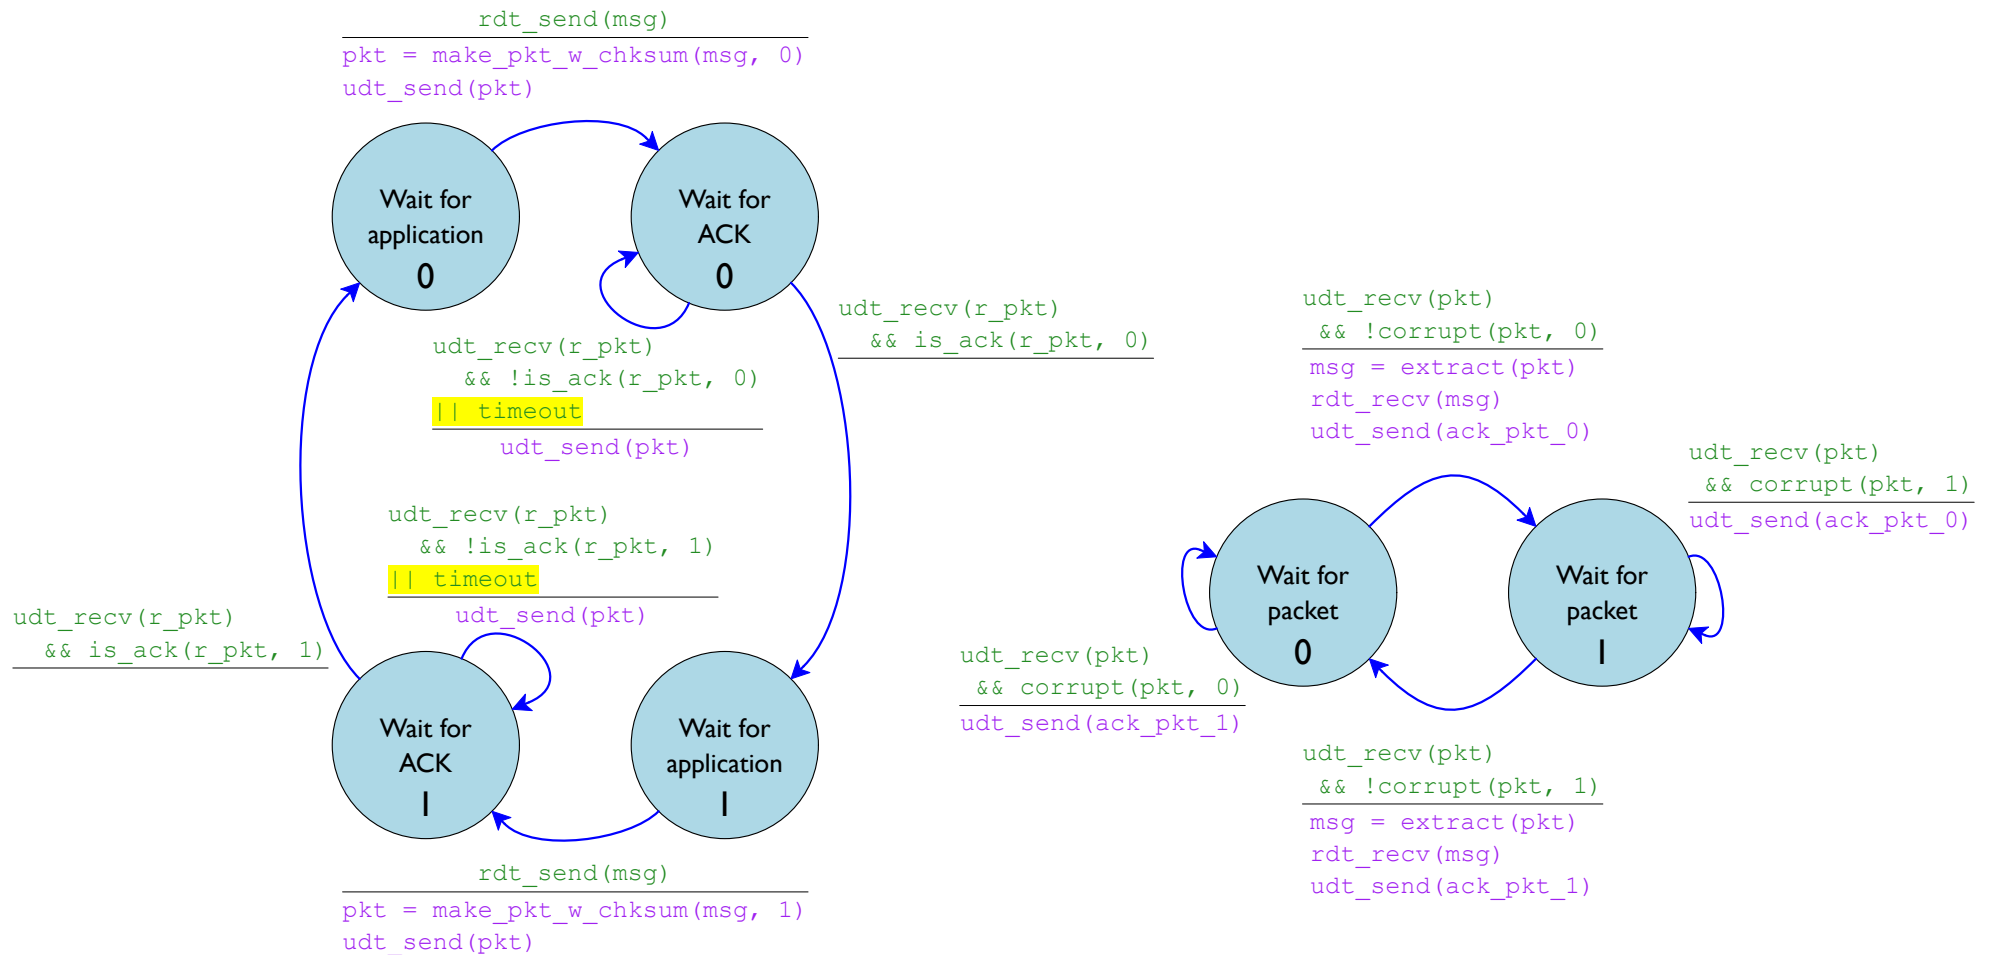
We can use `udt_send(ack_pkt_1)` as a NACK for `udt_send(ack_pkt_0)` and vice versa



Handling ACK Corruption

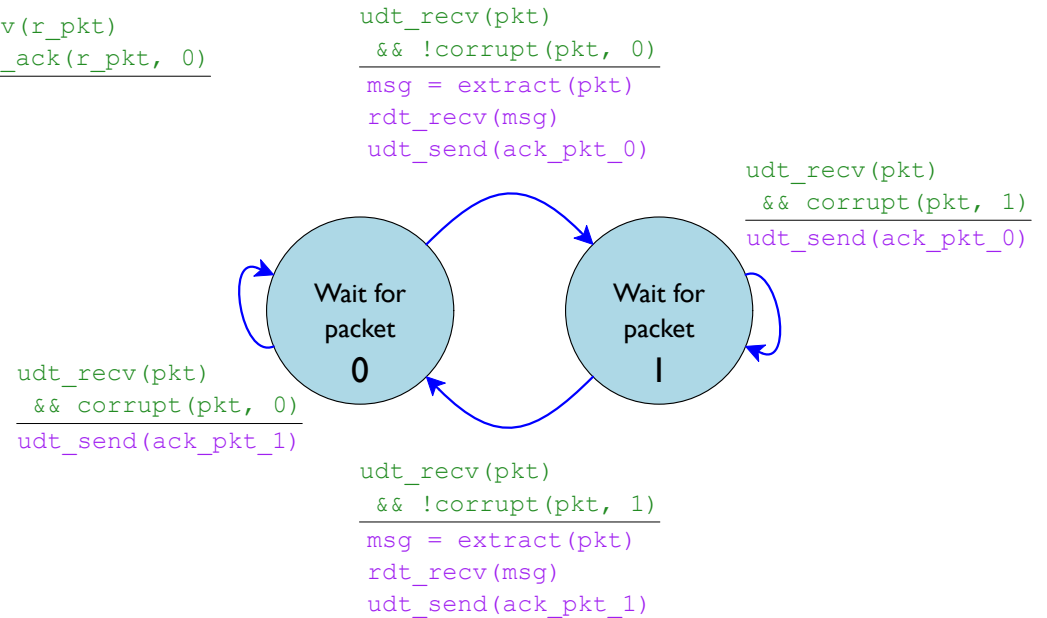
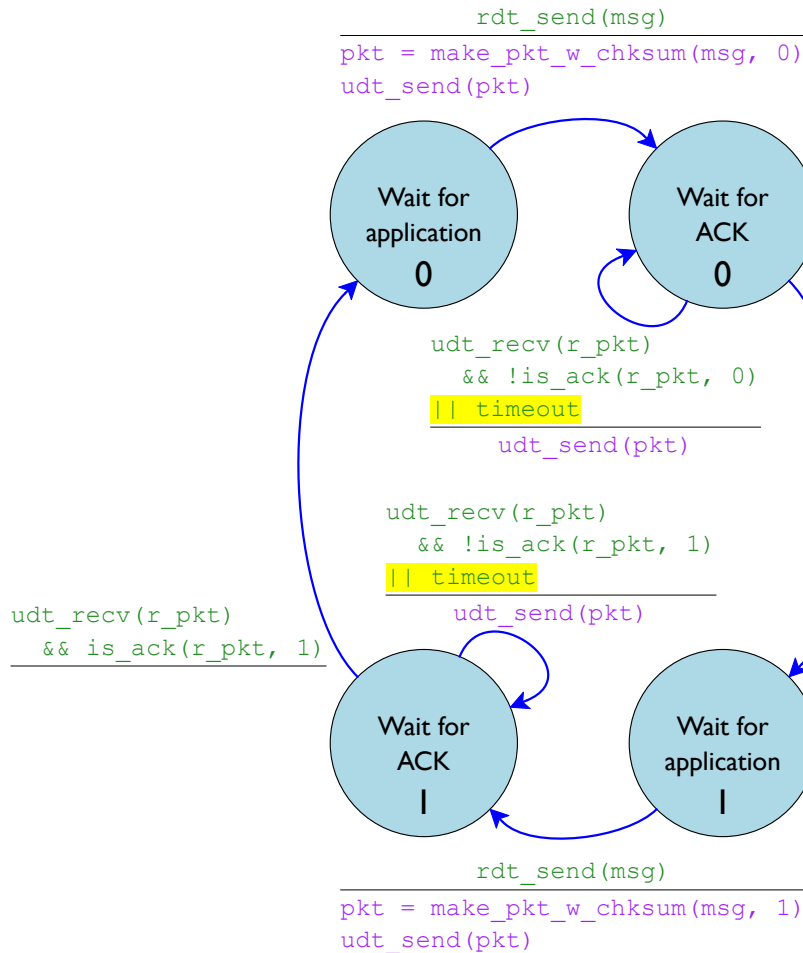


Handling Corrupt and Lost Packets



Handling Corrupt and Lost Packets

To handle duplicated and reordered packets, use a **sequence number** that always counts up instead of just 0 and 1




Choosing a Timeout

RTT is minimum useful timeout

- too small \Rightarrow resend data and ACKs unnecessarily
- too large \Rightarrow sender waits too long to resend


$scale \times avg(RTT) + stddev(RTT)$ is a good approach

Sequential Messages



Throughput is limited by latency


Pipelined Messages




Need a way to track multiple packets in flight

Buffers

sending host

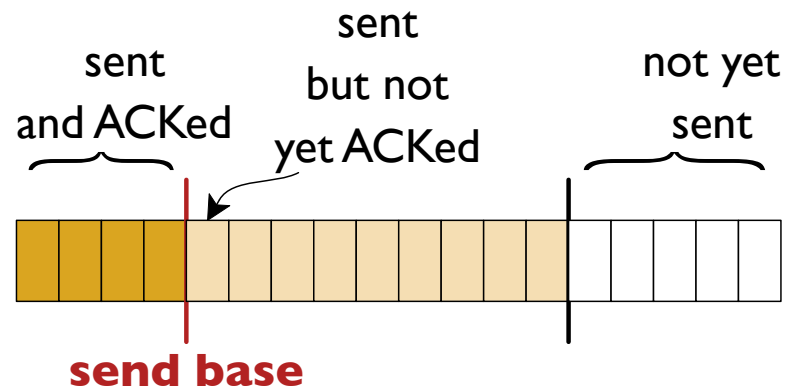


receiving host




Buffers

sending host

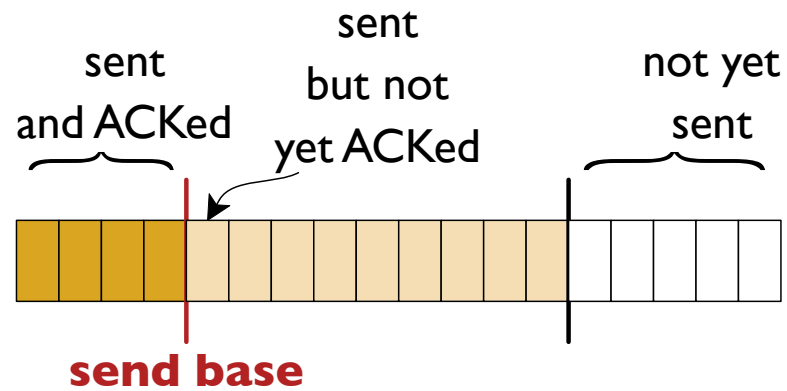


receiving host




Buffers


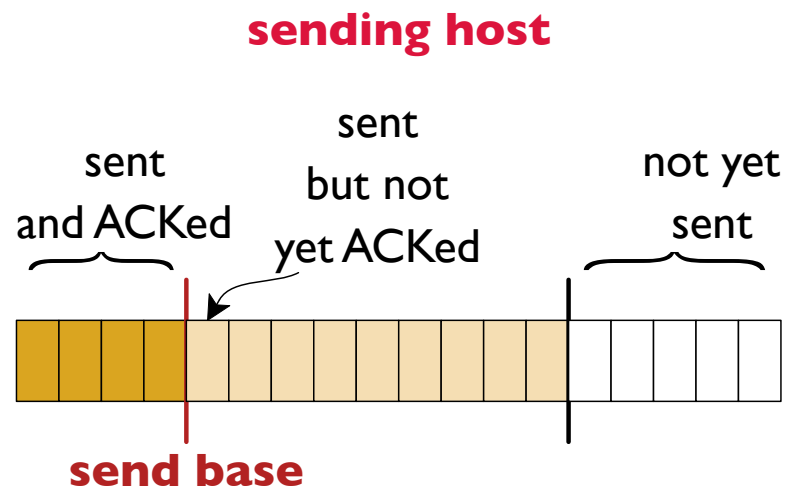
sending host




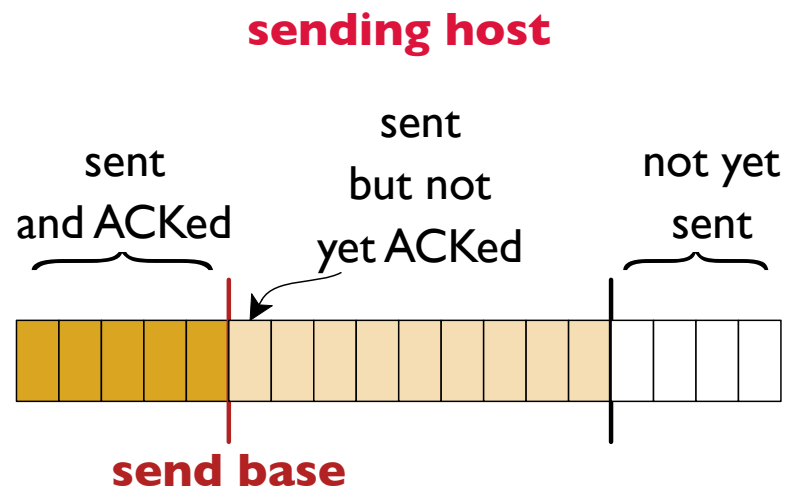
receiving host




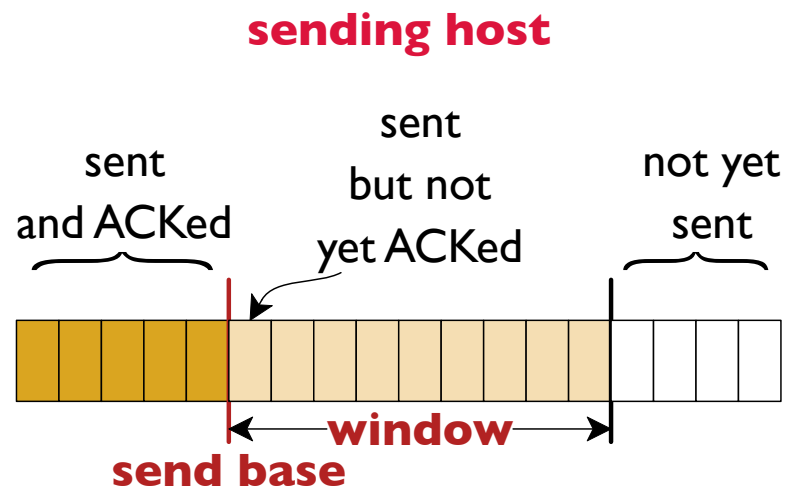
Buffers



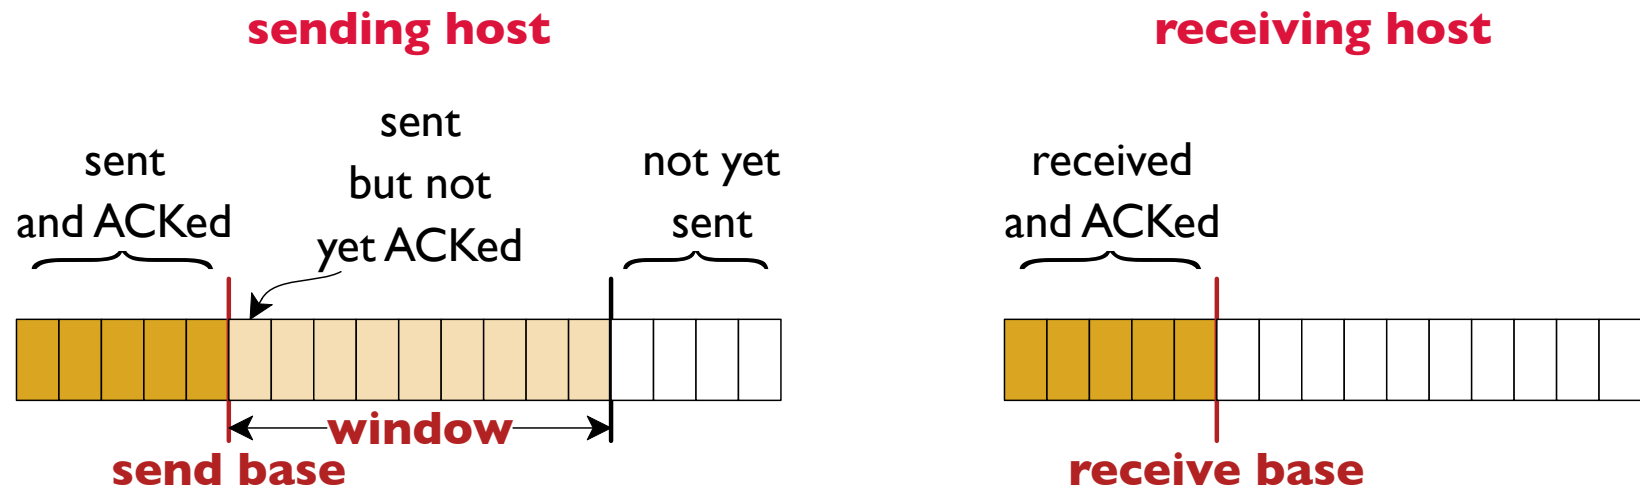
Buffers



Buffers

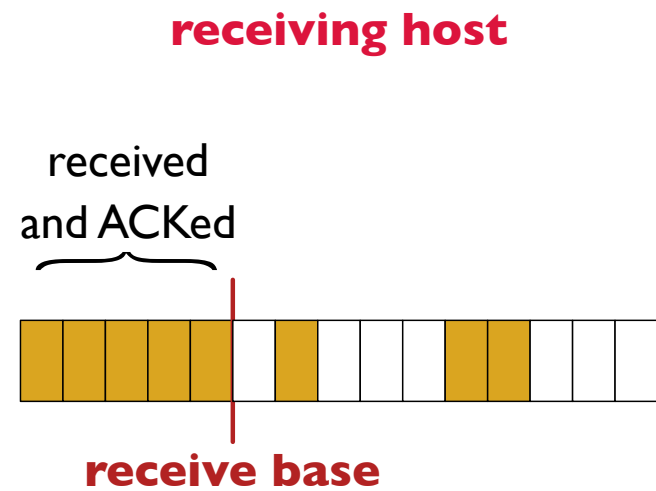
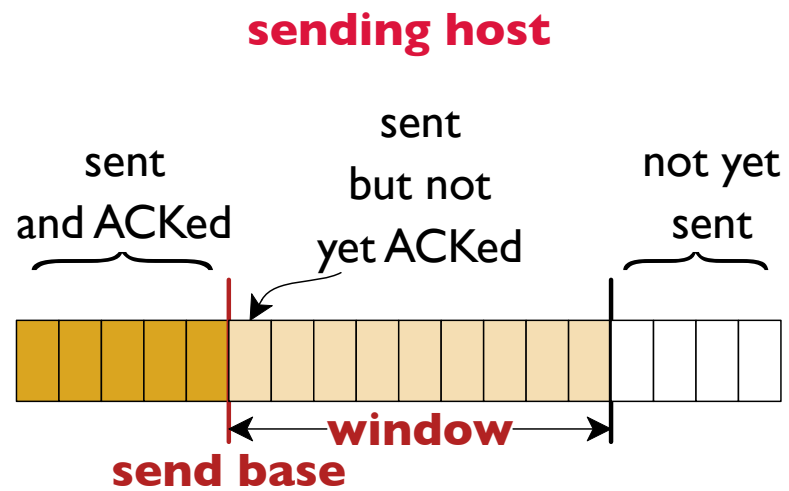


Buffers

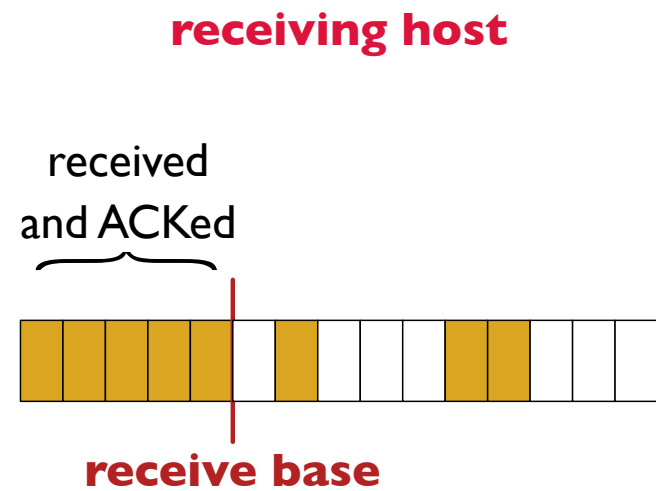
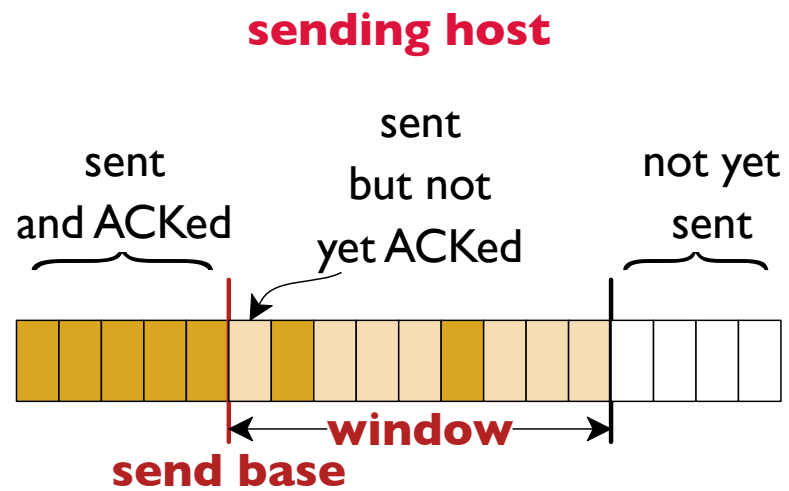


Like a timeout, the window size needs to be chosen well

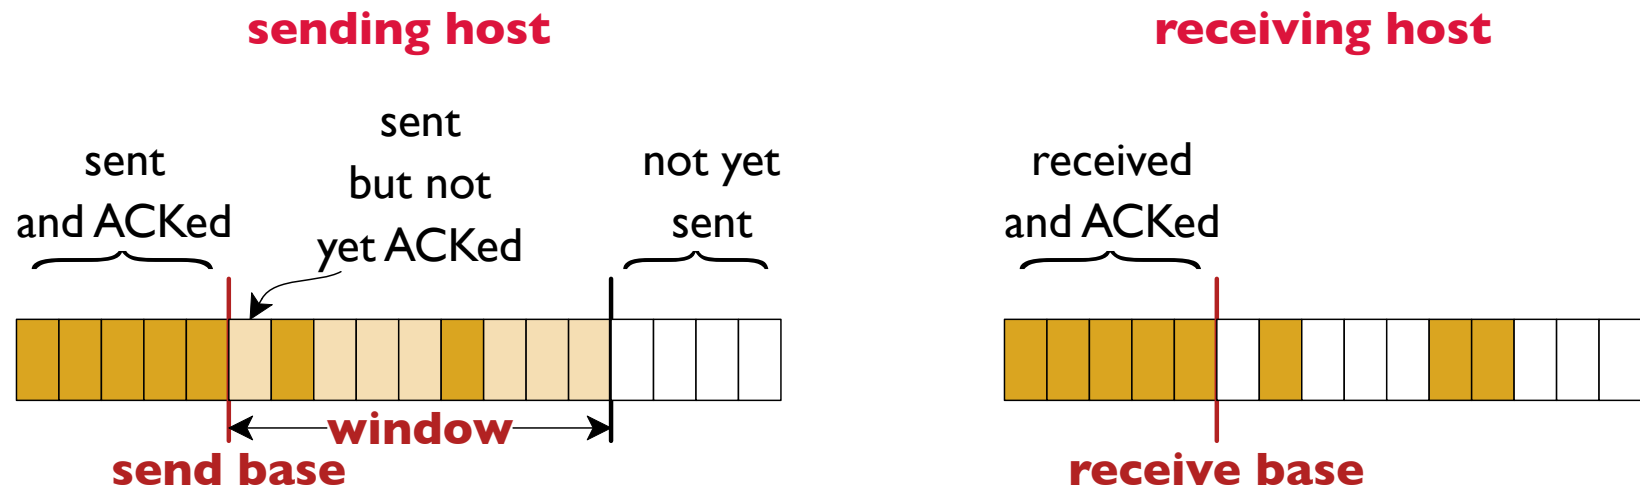
Buffers



Buffers

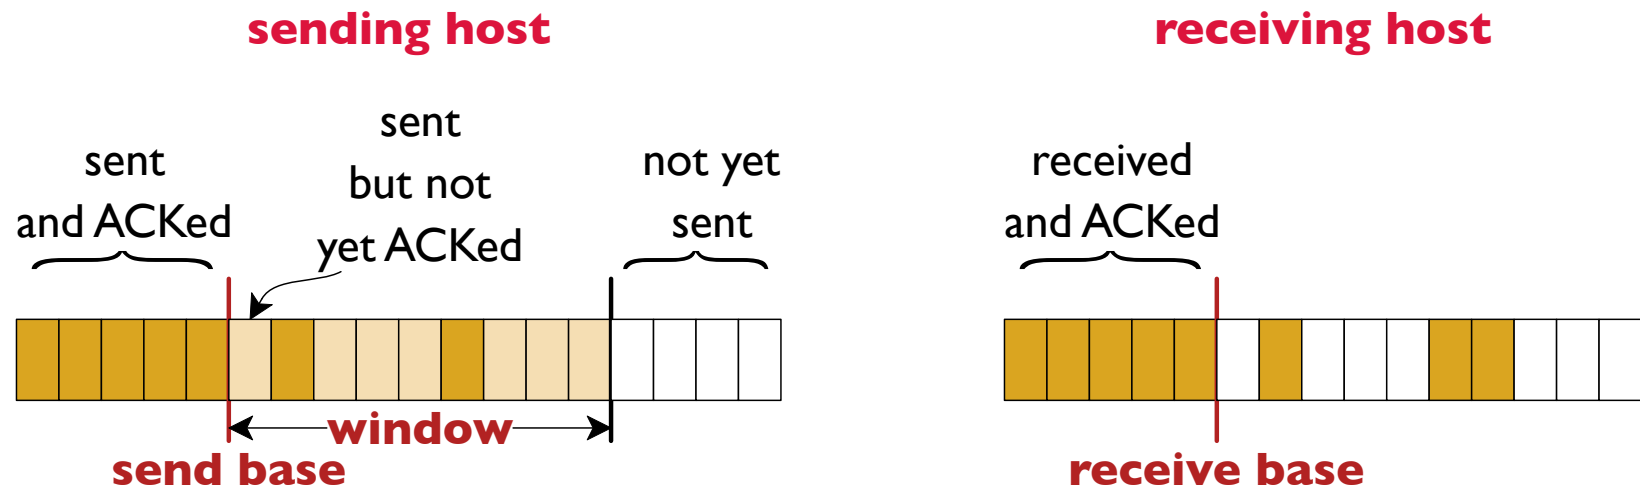


Buffers



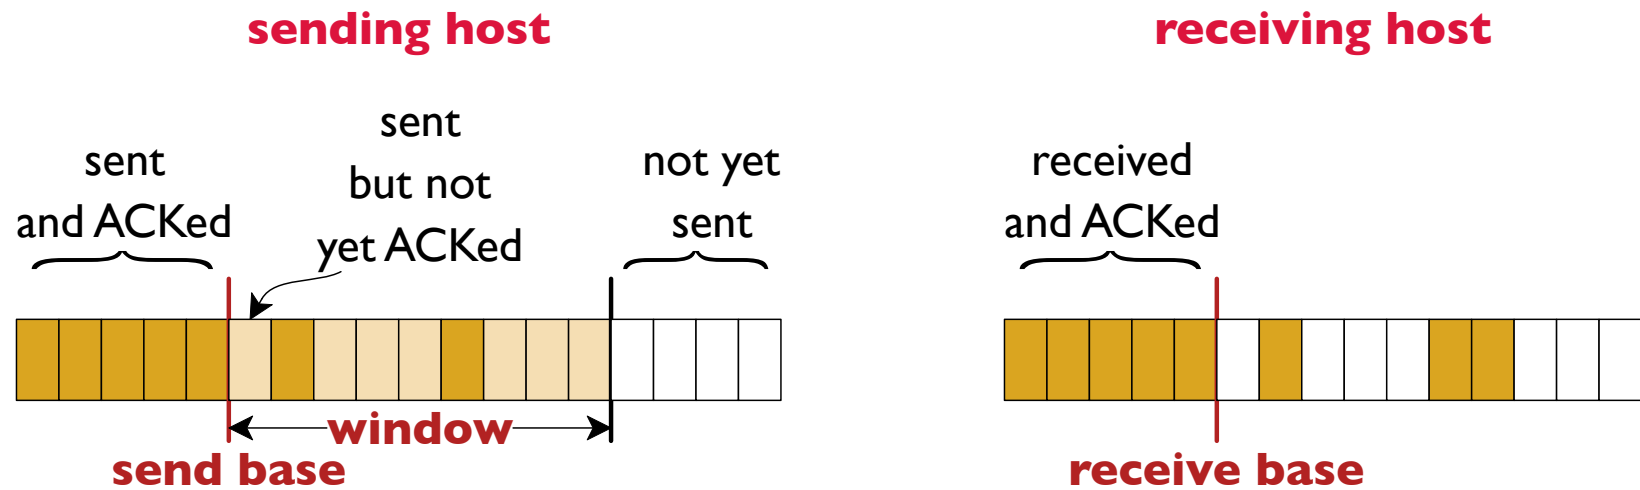
Selective repeat: on timeout, re-send unACKed

Buffers



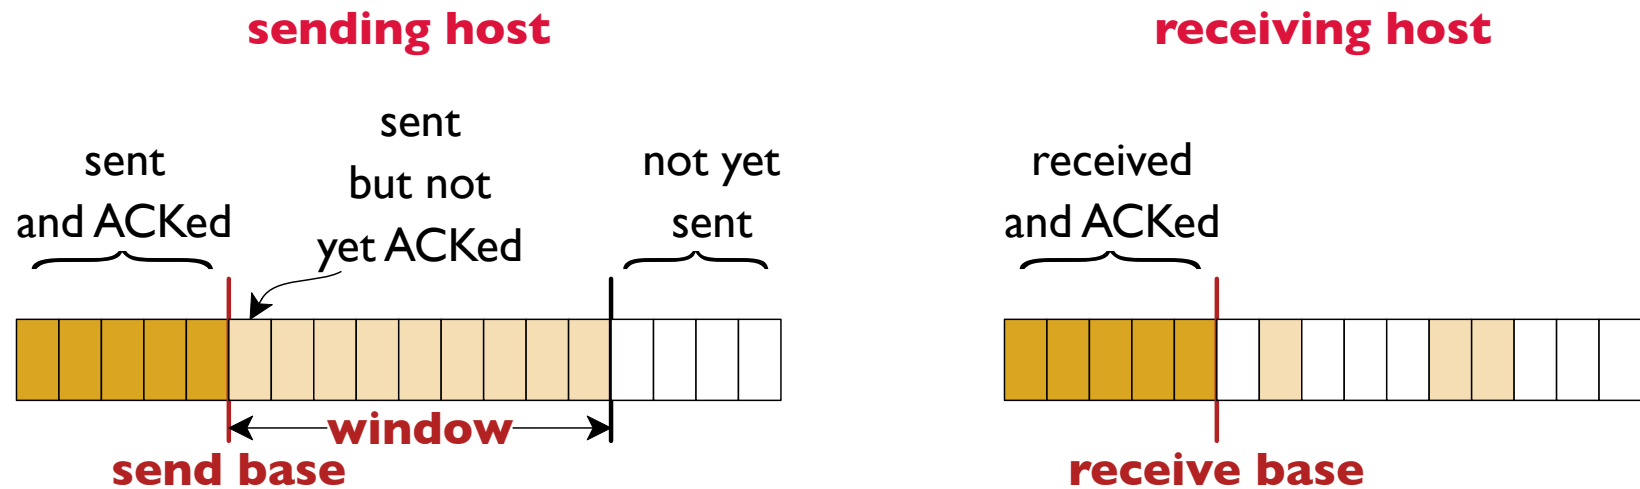
Selective repeat: on timeout, re-send unACKed
Each packet must be specifically ACKed

Buffers



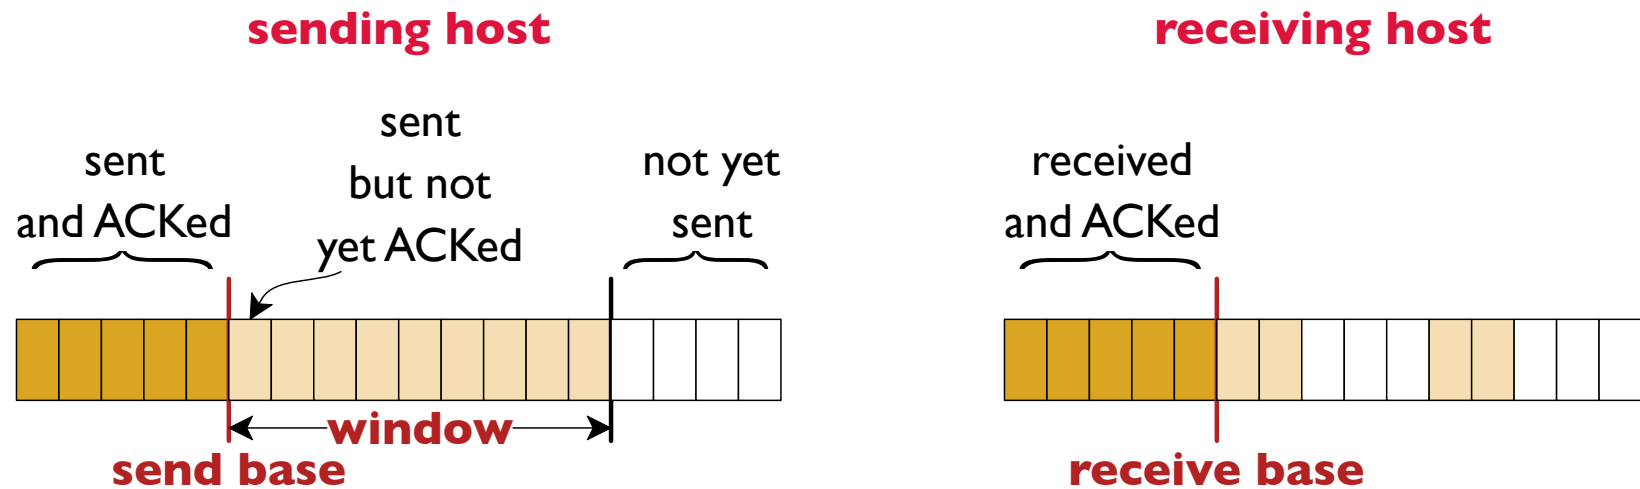
Go-Back-N: on timeout, re-send in window

Buffers



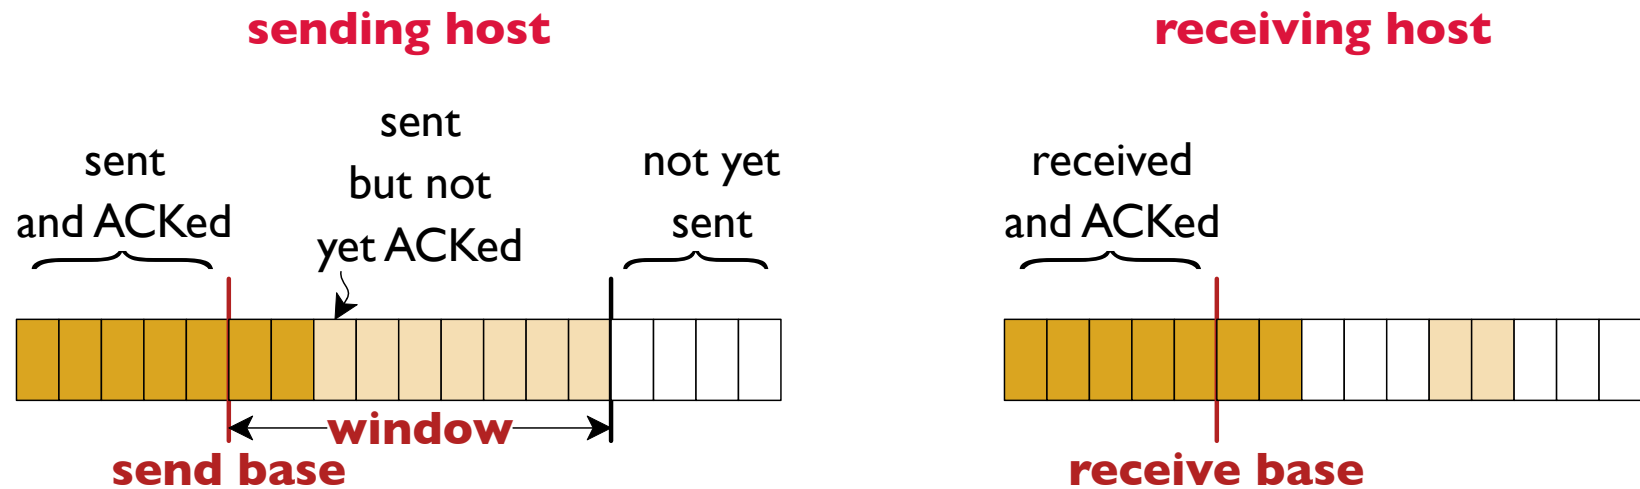
Go-Back-N: on timeout, re-send in window
Can use a **cumulative** ACK

Buffers



Go-Back-N: on timeout, re-send in window
Can use a **cumulative** ACK

Buffers



Go-Back-N: on timeout, re-send in window
Can use a **cumulative** ACK

Summary

Reliable data transfer can be implemented on top of an unreliable layer

State machines abstract over program details to explain just the program's states and transitions

- **ACKs** and **NACKs**
- **sequence numbers** for both ACKs and implicit NACKs
- **cumulative** ACKs versus **selective repeat**