Announcements

- Reading
  - Today: 6.3-6.4
  - Tuesday (after spring break): 6.5-6.6

Multiplexing in the Transport Layer

- Upward multiplexing
  - putting multiple transport connections onto one network connection
  - used to accommodate pricing strategies that charge for connections
- Downward multiplexing
  - using several network connections per transport connection
  - permits use of multiple copies of network resources
    - if the network layer uses sliding windows
      - a high latency network may under utilize the link
      - multiple connections each get a window
    - per connection buffer allocation
      - get more buffers
    - round-robin scheduling
      - get a larger share of link bandwidth

Crash Recovery

- Router or Link Crashes
  - Data in transit can be lost.
  - End nodes have sufficient state to recover lost data.
  - Transport protocol can hide network failures from the application.
- Host Crashes
  - Transport level state will be lost at one end.
  - Does the transport layer have sufficient info to recover?, **No!.**
    - Information must flow down to network and up to transport user
      - ACKs go down, and data goes up.
      - It is not possible to make these two operations atomic.
    - lack of stable storage causes this problem
  - Result, higher up layer must deal with host crashes

Protocol State Machines

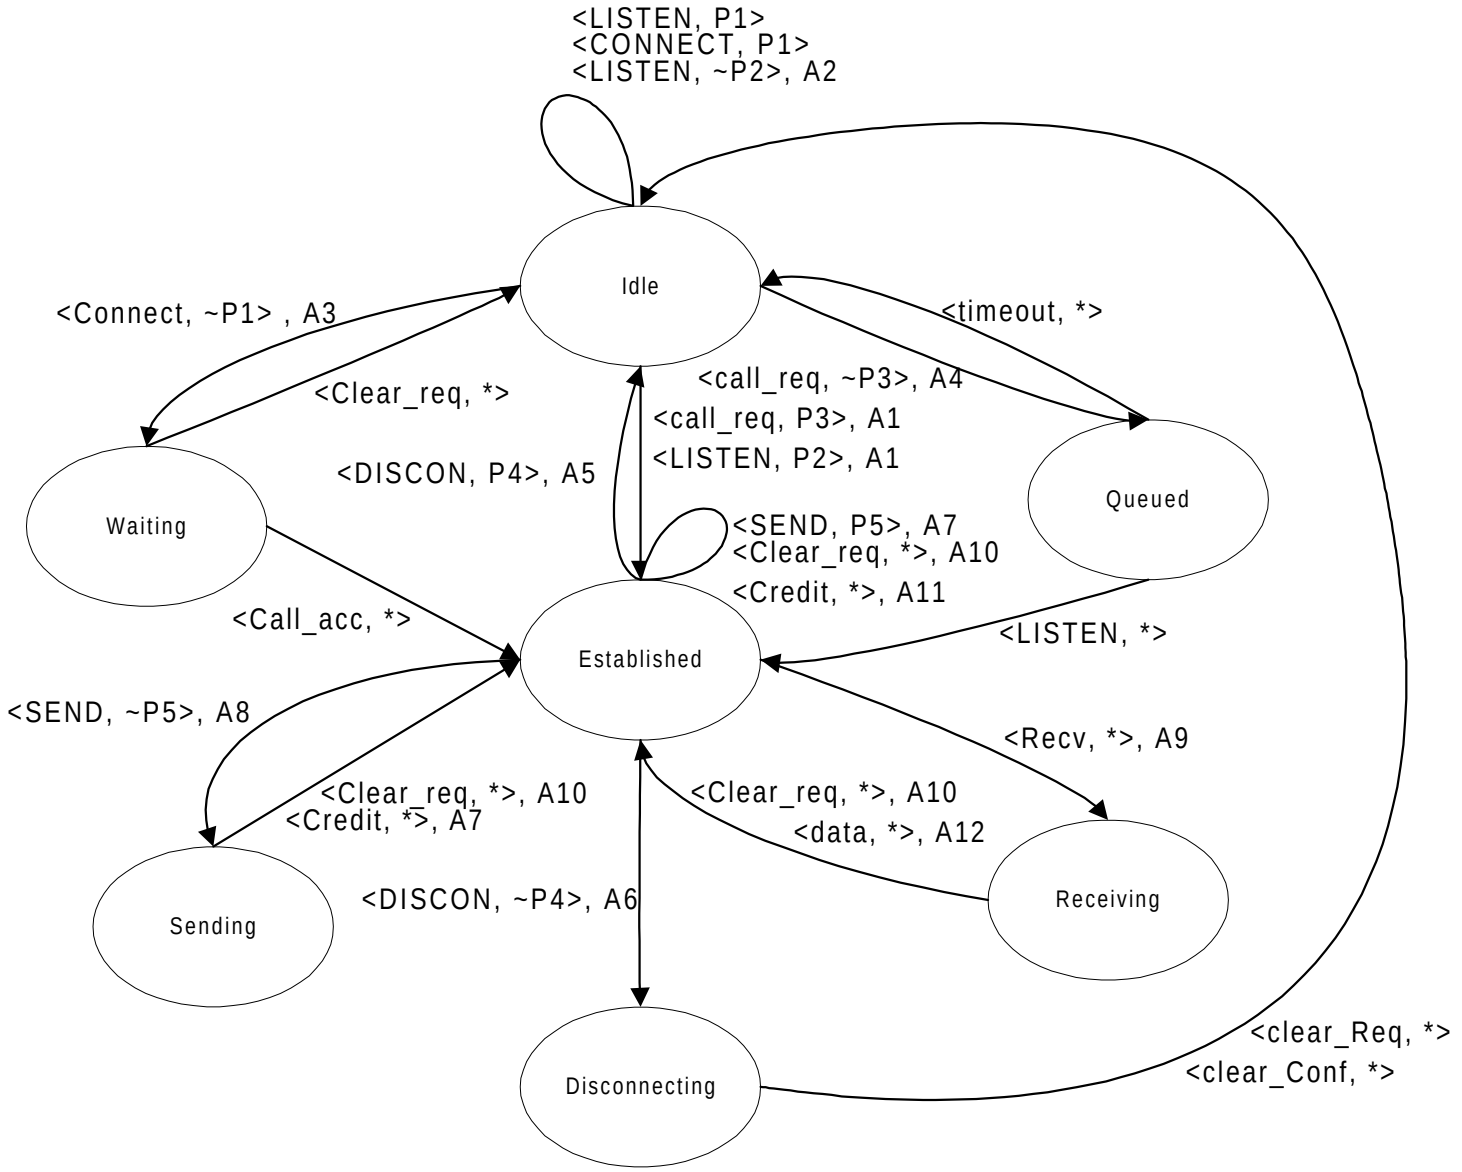

Predicates And State Transitions

Pred	Meaning
P1	Connection table full
P2	Call_req pending
P3	LISTEN Pending
P4	Clear_req Pending
P5	Credit Available

Act	Meaning
A1	Send Call_acc
A2	Wait for Call_req
A3	Send Call_req
A4	Start Timer
A5	Send Clear_conf
A6	Send Clear_req
A7	Send message
A8	Wait for credit
A9	Send Credit
A10	Set Clr_req_recv flag
A11	Record credit
A12	Accept message

TCP Protocol

- TSAPs
  - Use <host, port> combination
  - Well known ports provide services
    - first 256 ports
    - SMTP 25, Telnet 23, Ftp 21, HTTP 80
- Provides a **byte** stream
  - this is **not** a message stream
  - a message (single call to send) may be split, merged, etc.
- Urgent Data field
  - provides cut through delvery *within* a trasport connection
  - used to send breaks or other high priority info

TCP Packet Format

- Permits ACKs to be piggy packed
  - ACK is next byte expected
  - ACK is only valid if ACK bit is set
- Sequence number
  - first byte in packet
- Also used for connection establishment

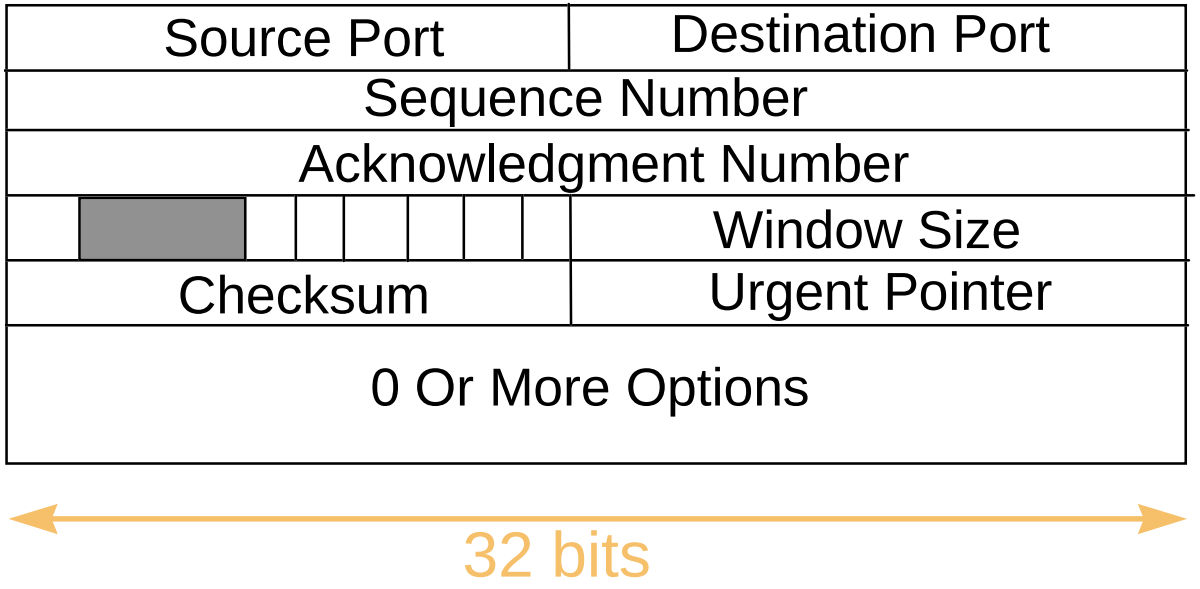

TCP Connection Management

- Three-way Handshake
- Initial Sequence Numbers
  - Use a 4 micro-second clock
  - hosts must wait T (120 seconds) before a reboot
- Connection Closure
  - Each side uses a FIN and FIN_ACK message
  - A FIN times out after 2 T (240 seconds)
  - Keep alives used to timeout half dead connections

TCP Flow Control

- Use Variable Sized Sliding Window
  - ACK indicates start of window
  - Window size indicates current size of window
- Receiver can send a window of 0
  - indicates that it want to pause connection
  - urgent data need not follow this request
- Window size of 16 bits is too small
  - 64K Bytes
  - only a small fraction of the in-flight bytes when
    - bandwidth is high
    - delay is high
  - solution: window shift option:
    - bit shift window up to 16 bits
    - permits up to 2^{32} byte windows
    - reduces window granularity