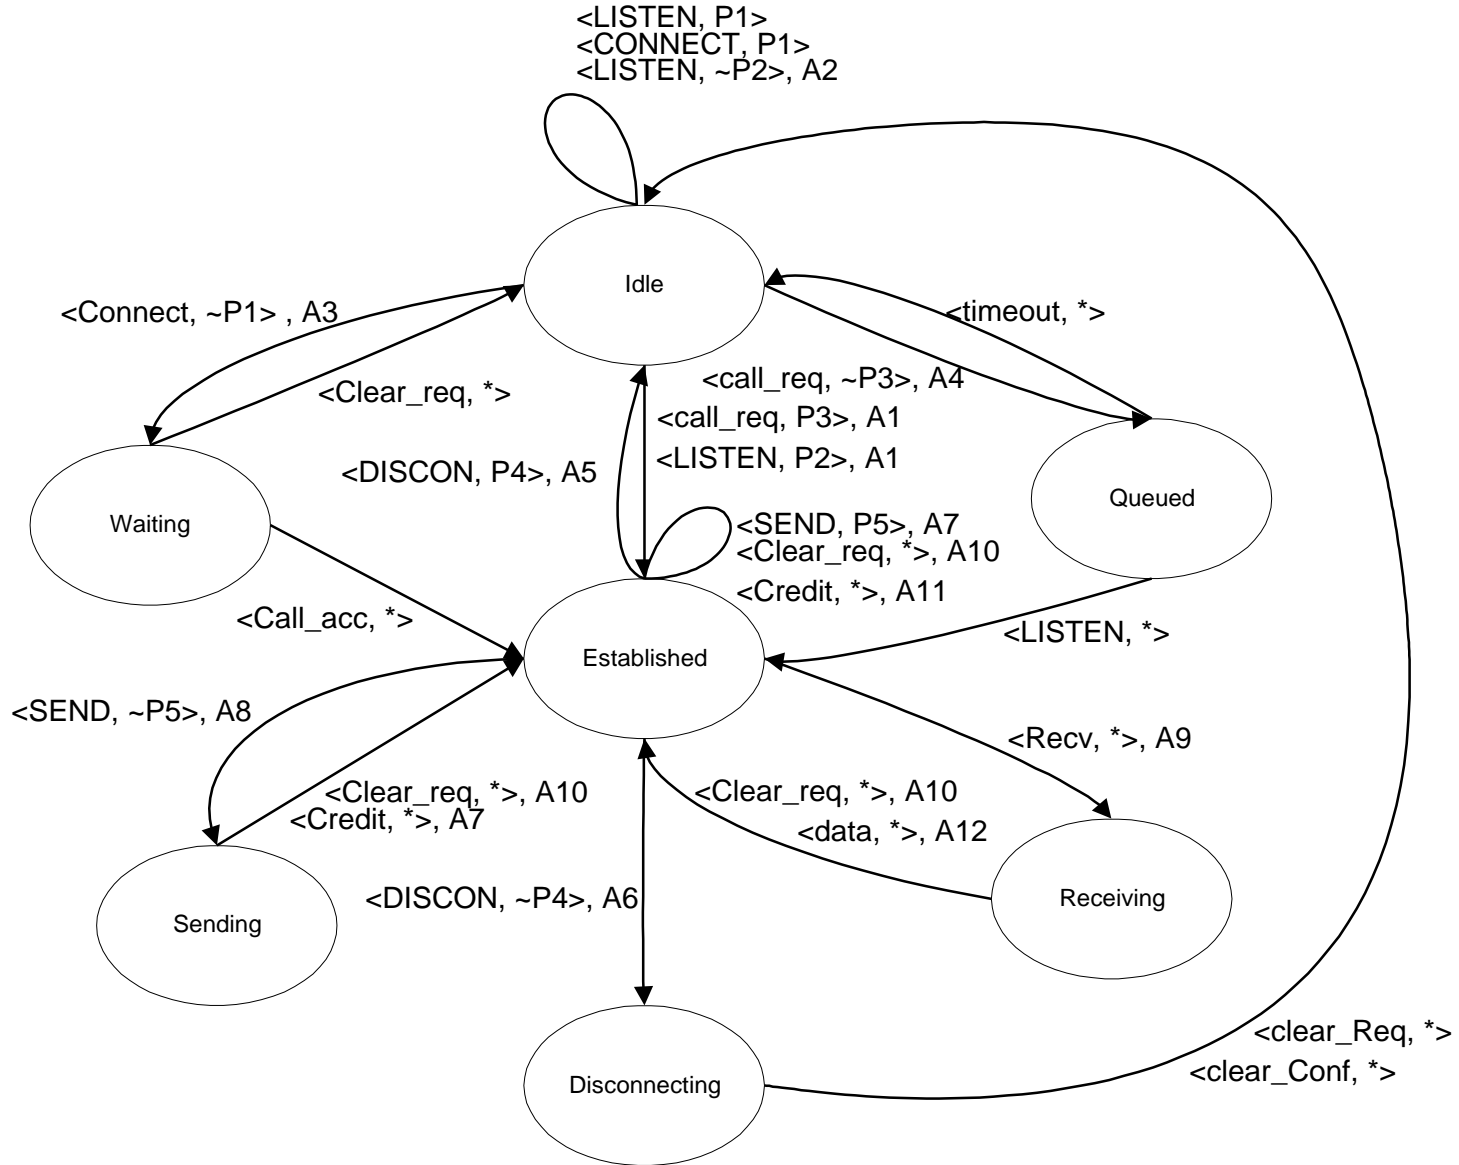


Announcements

- Reading
 - Today: Chapter 6 (6.3 and 6.4)
- Midterms: Requests for re-grade due by 10/19
- HW #1 (Due Today)
- Class Project Notes:
 - My comment about the parameter to `config_get_info` was **wrong**. It should take an array of 16 chars (I.e. a v6 address). See class web page for info about `parseV6Addr` which converts a string to a v6 address
 - the bit field notation `int x:y;` is **very** useful for packet headers

Protocol State Machines



Predicates And State Transitions

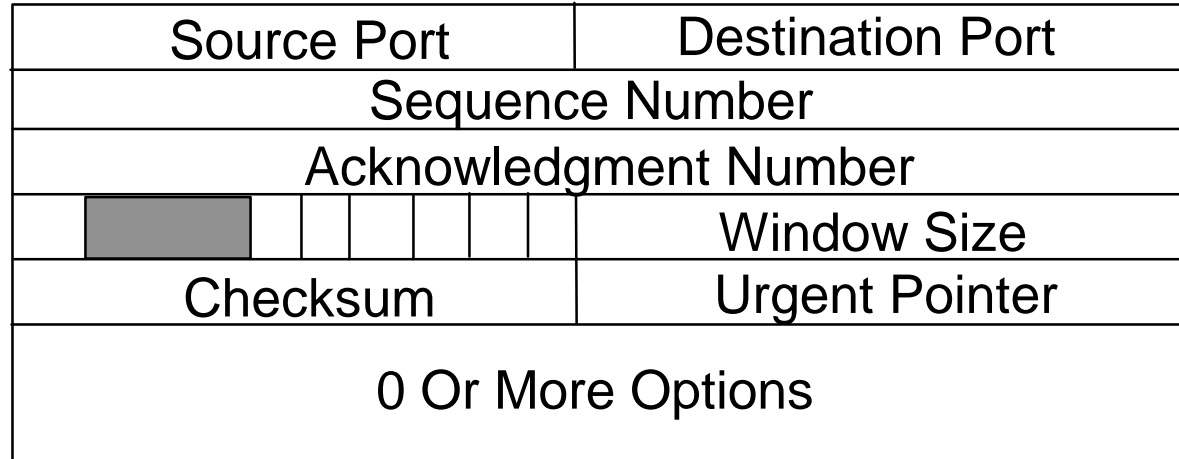
Pred	Meaning	Act	Meaning
P1	Connection table full	A1	Send Call_acc
P2	Call_req pending	A2	Wait for Call_req
P3	LISTEN Pending	A3	Send Call_req
P4	Clear_req Pending	A4	Start Timer
P5	Credit Available	A5	Send Clear_conf
		A6	Send Clear_req
		A7	Send message
		A8	Wait for credit
		A9	Send Credit
		A10	Set Clr_req_rcv 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 delivery *within* a transport 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



← 32 bits →

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 $2T$ (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

TCP Congestion Control

- Detecting Congestion

- In general it is difficult
- But, consider why a packet might be dropped
 - link error - but links are very reliable now
 - buffer overflow --> congestion
- Use re-transmission timeouts as an estimate of congestion

- Dealing with Congestion

- add a second window (congestion window)
 - limit transmissions to $\min(\text{recv window}, \text{congestion window})$
- start with congestion window = max segment window
 - initial max segment is one kilo-byte
 - on a ACK without a timeout
 - if $\text{window} < \text{threshold}$, increment by one max segment
 - otherwise increment by initial max segment
- on timeout
 - cut threshold in half
 - set window size to initial max segment