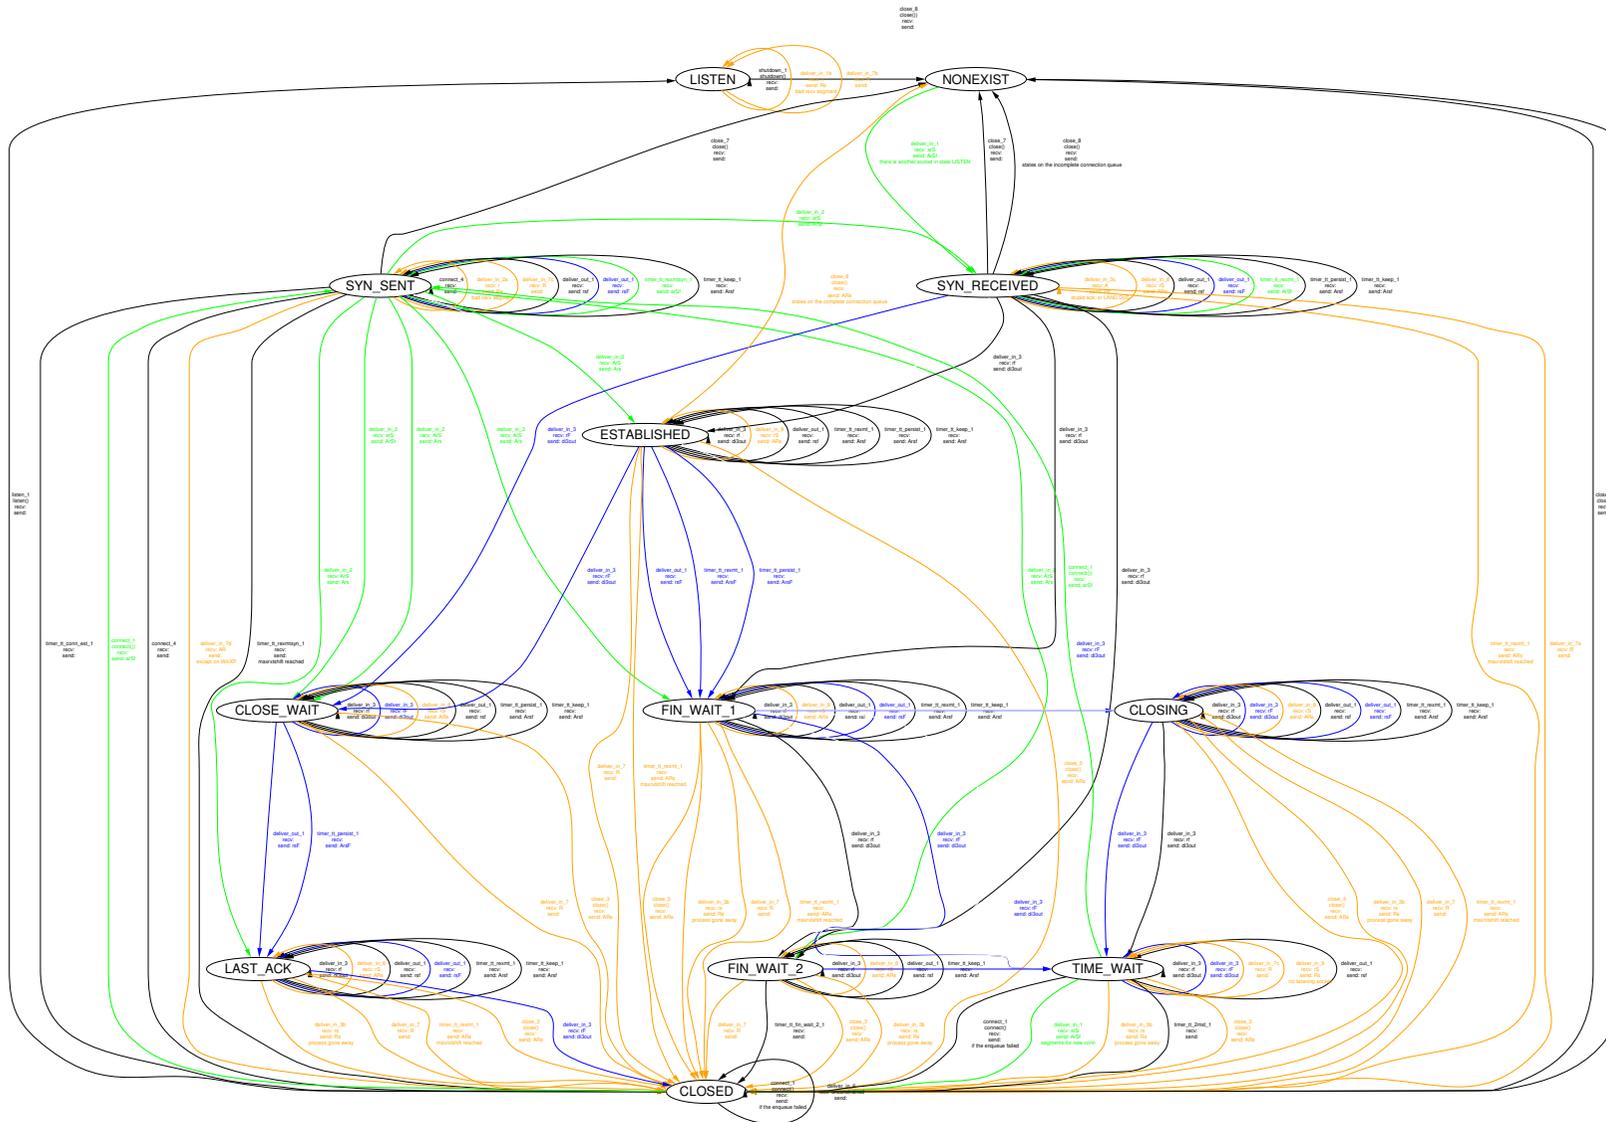


TCP: an approximation to the real state diagram



What Is This?

This graph shows an approximation to the Host Transition System of the TCP specification

TCP, UDP, and Sockets: rigorous and experimentally-validated behavioural specification. Volume 1: Overview. Volume 2: The Specification. Steven Bishop, Matthew Fairbairn, Michael Norrish, Peter Sewell, Michael Smith, and Keith Wansbrough. 2005.

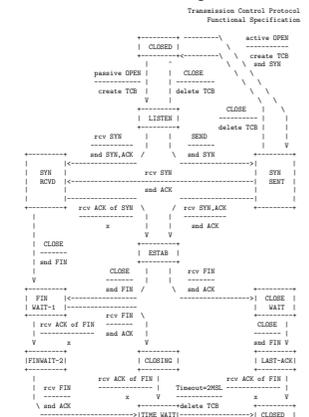
The states are the classic 'TCP states', though note that these are only a tiny part of the protocol endpoint state, in the specification or in implementations. The transitions are an over-approximation to the set of all the transitions in the model which (1) affect the TCP state of a socket, and/or (2) involve processing segments from the host's input queue or adding them to its output queue, except that transitions involving ICMPs are omitted, as are transitions arising from the pathological BSD behaviour in which arbitrary sockets can be moved to LISTEN states. Transitions are labelled by their Host LTS rule name (e.g. `socket_1_deliver_m_1`, etc.), any socket call involved (e.g. `close_1`), and constraints on the flags of any TCP segment received and sent, with e.g. `R` indicating that RST is set and `r` indicating RST is clear. Transitions involving segments (either inbound or outbound) with RST set are coloured orange; others that have SYN set are coloured green; others that have FIN set are coloured blue; others are coloured black. The FIN indication includes the case of FINs that are constructed by reassembly rather than appearing in a literal segment.

The graph is based on data extracted manually from the HOL specification. The data does not capture all the invariants of the model, so some depicted transitions may not be reachable in the model (or in practice). Similarly, the constraints on flags shown may be overly weak.

Transition Rules

- `close_1` Successful abortive close of a synchronized socket
- `close_2` Successfully close the last file descriptor for socket in the CLOSED, SYN_SENT or SYN_RECEIVED states
- `close_3` Successfully close the last file descriptor for a listening TCP socket
- `connect_1` Begin connection establishment by creating a SYN and trying to enqueue it on host's output queue
- `connect_2` Full socket has pending error
- `connect_3` Passive open: receive SYN, send SYN/ACK
- `connect_4` For a listening socket, receive and drop a bad datagram and either generate a RST segment or ignore it. Drop the incoming segment if the socket's queue of incomplete connections is full.
- `deliver_m_1` Completion of active open (in SYN_SENT) receive SYN/ACK and send ACK or simultaneous open (in SYN_SENT) receive SYN and send SYN/ACK
- `deliver_m_2` Receive bad or bogus datagram and RST or ignore for SYN_SENT socket
- `deliver_m_3` Receive data, FIN, and ACKs in a connected state
- `deliver_m_4` Receive data after process has gone away
- `deliver_m_5` Receive invalid ACK or LAMBO SYN in SYN_RECEIVED state
- `deliver_m_6` Receive and drop (silently) a same segment that matches a CLOSED socket
- `deliver_m_7` Receive RST and zap non-CLOSED, LISTEN, SYN_SENT, SYN_RECEIVED, TIME_WAIT sockets
- `deliver_m_8` Receive RST and zap SYN_RECEIVED socket
- `deliver_m_9` Receive RST and ignore for LISTEN socket
- `deliver_m_10` Receive RST and ignore for SYN_SENT (unacceptable ack) or TIME_WAIT socket
- `deliver_m_11` Receive RST and zap SYN_SENT (acceptable ack) socket
- `deliver_m_12` Receive SYN in non-CLOSED, LISTEN, SYN_SENT, SYN_RECEIVED, TIME_WAIT state
- `deliver_m_13` Receive SYN in TIME_WAIT state if there is no matching LISTEN socket or sequence number has not increased
- `listen_1` Common case TCP socket
- `listen_2` Successfully put socket in LISTEN state
- `listen_3` Successfully reset socket in the LISTEN state from any non-CLOSED, LISTEN) state on FreeBSD
- `shutdown_1` Shut down read or write half of TCP connection
- `socket_1` Successfully create a new file descriptor for a fresh socket
- `timer_1` 2MSL timer expires
- `timer_2` connection establishment timer expires
- `timer_3` FIN_WAIT_1 timer expires
- `timer_4` lingering timer expires
- `timer_5` persist timer expires
- `timer_6` retransmit timer expires
- `timer_7` SYN retransmit timer expires

The RFC793 Original



<http://www.cl.cam.ac.uk/users/pes20/Netsem>
March 18, 2005