


Quick-Start for DCCP


draft-fairhurst-tsvwg-dccp-qs-00
(Individual Submission)

Gorry Fairhurst
Arjuna Sathiaseelan

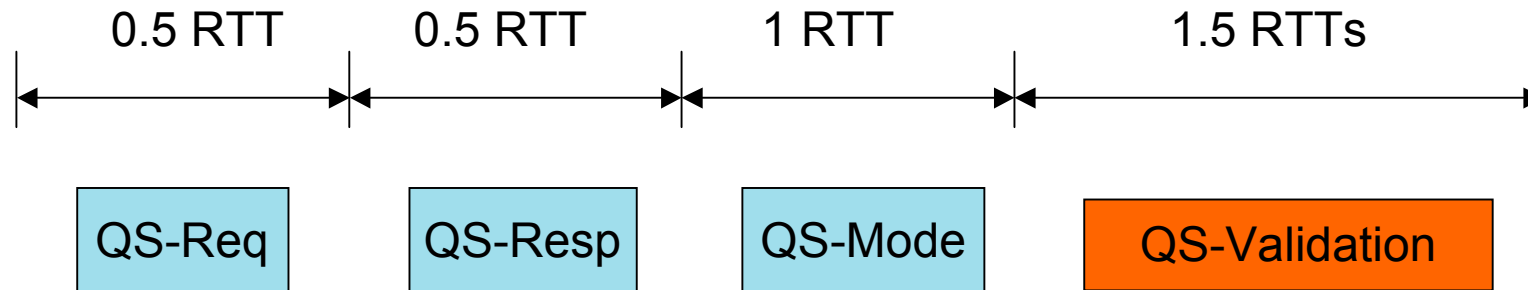


- Similar to QS with TCP [RFC 4782].
- Sender MAY use a Quick-Start request:
 - *At start of a connection.*
 - *In the middle of a connection.*
- SHOULD send the request on a packet that requires an acknowledgement (DCCP-Request, DCCP-Response, or DCCP-Data).
- MUST NOT make a subsequent Quick-Start Request within four RTTs.
 - *CCID-3 responds slowly to changes.*

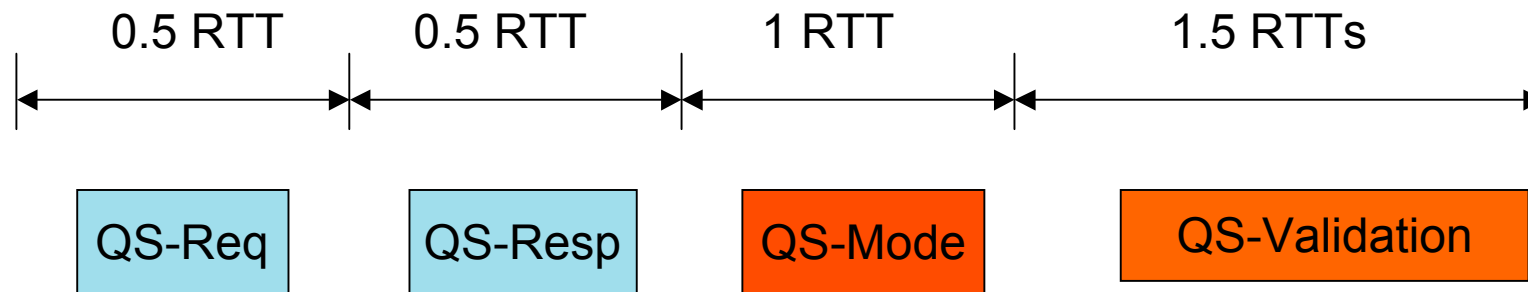
- QS-Requests processed identical to TCP.
- On receipt of QS-Request:
 - Receiver **SHOULD** immediately send a QS-Response.
- Receiver aligns feedback to end of period with Quick-Start Packets.



- Sender SHOULD NOT ignore a feedback packet with QS-Response option.
- Sender enters QS-Mode.
 - Sending host sets Quick-Start sending rate
 $QS\text{-sendrate} = R * s / (s + H)$
 - CCID 3 is rate paced protocol.
QS packets are naturally rate paced.
- Sender exits QS Mode when either:
 - A feedback packet acknowledging one or more Quick-Start packets.
 - Detection of a loss event.
 - A period of one RTT after receipt of QS-Response.
- If no reported loss (or ECN marking), enters Quick-Start Validation Phase




- Unlike TCP, TFRC receives a feedback once per RTT.
- Add a new Quick-Start Validation Phase.
 - Period to affirm capacity used by QS packets did not introduce congestion.
 - Sender tentatively permitted to continue sending at QS-sendrate.
 - Limited to a maximum of 1.5 RTTs (or loss, or feedback for QS Packets)
- At the end of the Quick-Start Validation phase:
 - Sender stops using the QS-sendrate.
 - Uses the standard congestion control mechanisms.



- If no feedback received within Quick-Start Mode or Validation Phase:
 - **MUST** return to minimum of original rate (at start of QS Mode) and one half of QS-sendrate.
- If a feedback packet arrives reporting packet loss
 - **MUST** immediately leave the Quick-Start Mode or Validation Phase.
 - Enters congestion avoidance phase.
 - Re-calculate send rate X :

$$X = \max(\min(X_{\text{calc}}, \min(2 \cdot X_{\text{recv}}, 2 \cdot \text{QS_recv-rate})), s/t_{\text{mbi}}).$$



- After idle, must requested QS rate consider current loss event rate?
 - Is it appropriate use QS after receiving a loss-event, as a way of getting more capacity than allowed by the throughput equation ?
- Need to consider implications of alternate paths, etc
 - Examine if there are specific issues
- Need to consider implications of over-shoot
 - QS approved capacity was not actually available, and DCCP sends at a higher rate for several (≤ 1.5 extra) RTTs.
- Other issues? Simulation...

- CCID-3 feedback timer
 - Receiver can use window counter not feedback timer expiry.
- Add CCID-2 text (should be easy)
- Which WG?
 - TSVWG or DCCP?