DCCP:
Issues From the Mailing List

Sally Floyd, Eddie Kohler,
Mark Handley, et al.

DCCP WG
March 4, 2004

Quotes from the Mailing List:

- "Real Time apps have no interest in being 'fair'."
- "Don't expect people like XXX to be happy with the 'you must fit in to TCPs view of the world' when most of their real-time applications are already being good citizens (by not sending 'all they can' when they don't have to)."
  - [Referring to a silence-suppressed, CBR, 4 Mbps video stream!]
- "Maybe some codecs of today, which were designed assuming a QoS enabled network, cannot make use of TRFC or DCCP."
- "TCP is broken. Nowhere is it written that TCP=best effort."

Background assumptions (ours):

- DCCP and best-effort traffic:
  - DCCP with best-effort service does not necessarily meet the needs of all apps.
  - In fact, best-effort service does not necessarily meet the needs of all apps.
  - DCCP is intended to solve the best-effort problem, not the QoS problem.
  - We believe that in the long run DCCP offers better performance than UDP to applications (e.g., ECN, NAT traversal, etc.)

Issue:
Steady-state fairness with TCP?

- RFC 2914, Congestion Control Principles, September 2000, BCP.
  - Preventing congestion collapse.
  - Sharing bandwidth reasonably fairly with TCP.
- The IETF doesn't control what is deployed in the Internet.
  - The IETF controls what is standardized in the IETF.
  - The current IETF will not standardize transport protocols for best-effort service that do not have adequate end-to-end congestion control.

Issue: Slow-Start

- CBR app writers don't want to slow start
  - ...after idle periods
  - ...ever
- CBR app writer perceptions [NB not direct quotes]
  - "We're sending at a low rate so why bother?"
  - Idle periods: "We're benefitting the network by going idle, why penalize us by forcing slow start after a quiet period?"
  - "Our traffic is more financially valuable to ISPs so congestion rules don't apply”
  - "TCP must be fixed [to be friendlier to CBR apps]"

Issue: Slow-Start

- CCID3 specifies initial rates of 4 pkts/RTT.
  - Recommends investigating initial rates of 8 ptks/RTT for small packets.
- For CBR apps with higher rates, this means that some initial packets could be ‘dropped’ by DCCP.
- Best-effort traffic with higher initial rates?
  - My own view:

    Explicit feedback from routers is needed.
  - E.g., Quick-Start, expired draft

    draft-amit-quick-start-02.txt.
  - You could help make this happen!

Issue: Limitation of at most doubling the sending rate

- Thread triggered by earlier user guide suggestion:
  - Send 2x your nominal rate to avoid:
    - getting penalized by "greedy" TCPs
    - slow start after idle periods
- But TFRC isn't penalized by TCP flows:
  - Transmit rate limited by *loss rate* not current rate.
- The limitation of at most doubling the sending rate remains (above a minimum rate):
  - A problem for bursty apps, instant-on apps, silence suppression.
  - When can this limitation be safely relaxed?

Issue: Slow-Start after Idle

- Proposal: Faster Start
  - Initial rate 8 pkts/RTT (instead of 4).
  - Quadruple rate each RTT up to previous rate (instead of doubling).
  - Until a drop or mark.
  - This needs further investigation.
- Implementation experience about slow-start problems will help.

Issues: apps with fixed rates, or a small number of possible rates, or limited to downshifting.

- Email: For some apps, users prefer fixed rates.
- DCCP can be used by fixed-rate apps.
  - Modulo slow-start, restart-after-idle issues.
  - DCCP will send at a sending rate allowed by the overall packet drop rate.
  - As always, implementation experience is needed.
- Proposal: for the apps above, DCCP could sometimes send as much as twice the "allowed" sending rate?
  - This requires a new CCID, and some further work.

New viewgraphs:

Issues: CBR flows

- Advice for CBR flows:
  - Monitor the steady-state packet drop rate, stop sending when the drop rate is too high.
  - "IAB Concerns Regarding Congestion Control for Voice Traffic in the Internet”, approved as an RFC.

Issues: Special Consideration for CBR Traffic?

- What if all of the traffic is CBR?
  - What about the two hours after an earthquake?
  - What about a chronically-congested link?
- What about the congested link where the TCP traffic is backing off, and the CBR traffic is causing the high drop rate?
- The TCP traffic is not all bulk-data transfers:
  - E.g., my web traffic making plane reservations.