Low Extra Delay Background Transport

draft-shalunov-ledbat-congestion-00

Stanislav Shalunov ⟨shalunov@bittorrent.com⟩

IETF 75, Stockholm
LEDBAT WG
July 29, 2009

Goals reminder

- saturate the bottleneck
- keep delay low with no other traffic
- yield to TCP
- add little extra delay

Receiver reminder

on data:

create ack
ack.ackseq = ...
ack.delay = time() - data.time
send ack

Rough pseudocode reminder (sender)

on ack:

```
current_delay = ack.delay
base_delay = min(base_delay, current_delay)
queuing_delay = current_delay - base_delay
off_target = TARGET - queuing_delay
cwnd += GAIN * off_target/cwnd
```

on loss:

```
cwnd /= 2
```

(see the draft for more precision)

Framing

- General-purpose congestion control loop
- Works over UDP
- Possibly a TCP mod/option in the future
  - need timestamps
- Possibly a DCCP CCID
- Possibly an SCTP mod/option

No late-comer's advantage

- Late-comer's advantage is based on incorrect base measurements by a late comer
- Enough dips with realistic numbers of flows to get a real base measurement every once in a while

No multiplying target by number of flows

- All flows measure the same queue
- All flows target the same queue size
- That's the size you get, no multiplying

Fairness by random redistribution

- What if we added a small random quantity to all measurements?
  - It's there already, no need to add
  - Flows see slightly different queue size
  - Depends on, e.g., your arrival relative to serialization of previous packet
  - Also small errors introduced by scheduling

Fairness by random redistribution, cont.

- With an extra deviate in measured delay, flows give or take a little on each packet
- Total used capacity stays the same on RTT timescale (absent change in cross-traffic)
- Random redistribution
- Rounds of random redistribution → comparable capacity for each flow

Parameter values

- Target of 25ms? 5ms? 50ms? 3ms?
- (Gain of 1 MSS/RTT? 10?)
- The values in the draft work
- A wide range of other values works, too

Target choice considerations

- Higher values are more robust
- Lower values add less delay
- At the low end, diminishing returns
- Human perception threshold a useful benchmark

QUESTIONS?