Performance Measurements of QUIC Communications

Algorithm to improve connection RTT evaluation using 1 bit more

Authors:
Fabio Bulgarella (Telecom Italia - TIM)
Mauro Cociglio (Telecom Italia - TIM)
Giuseppe Fioccola (Huawei Technologies)
Guido Marchetto (Politecnico di Torino)
Riccardo Sisto (Politecnico di Torino)

Montreal ANRW ‘19

Purpose of the work

- We are proposing an alternative way of improving the spin bit performance in delay measurements.
- It uses the delay bit which uses only one additional bit instead of the two required by the VEC (Valid Edge Counter).
- In this way, not all the three reserved bits are used, but one remains free for other purposes, such as for example, measurement of loss rate.

The latency Spin Bit

- The spin bit is a simple mechanism which causes one bit in the header to 'spin', generating one edge (a transition from 0 to 1 or from 1 to 0) once per end-to-end RTT.
- On-path observers can measure the time elapsed between these edges to generate one RTT sample per flow per round-trip period.
- **SERVER REFLECTS**: it sets the spin bit of outgoing packets to the same spin bit of the last packet received.
- **CLIENT INVERTS**: it sets the spin bit of outgoing packets to the opposite spin bit of the last packet received.

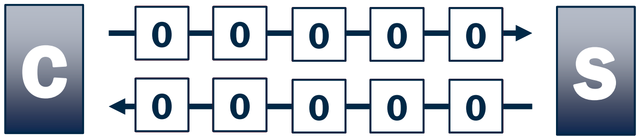

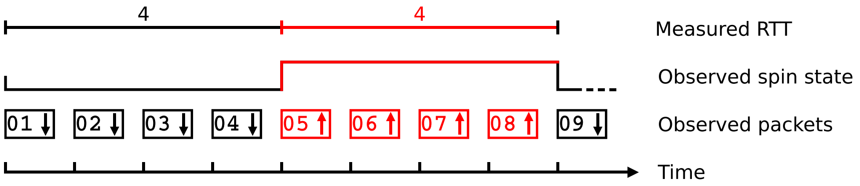

Spin Bit limitations

- Packet loss will tend to cause wrong estimates of RTT due to period width changes.
- Reordering of a spin edge will cause drastic underestimates of RTT since it will cause multiple edges to be observed per RTT.

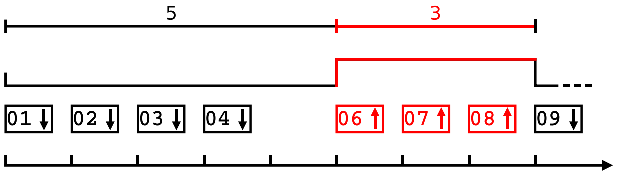

Edge loss

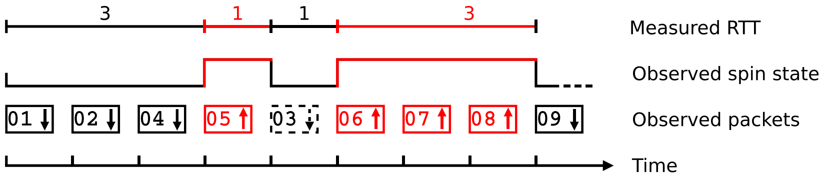

Edge reordering

- Application-limited sender can introduce delay in the edge reflection.
- Issues addressed by the VEC, a two-bit validation signal used to mark valid edges

The Valid Edge Counter method

- The VEC is a two-bit signal added to each packet whose purpose is to explicitly report whether an edge was valid when transmitted by the endpoint.
- A value greater than zero is assigned exclusively to valid edges. Then, when an endpoint detects an incoming packet carrying a spin transition, the VEC value of the next generated edge is set to the value contained in the received packet incremented by 1 (holding at 3).
  - Basically, the value of the VEC is increased every time a valid edge is reflected by one of the two endpoints, counting the number of semi-paths correctly crossed by the edge without incurring network impairments.
  - Instead, when the endpoint detects an impairment such as a reordered or lost edge, the VEC is set back to 1 so that the observer avoids completing incorrect measurements.

The Delay Bit

- The idea is to have a single packet, with a second marked bit, called «delay bit», that bounces between client and server. This packet is also called ‘’Delay Sample’’.
- A passive observer, placed on whatever direction, can compute the differerence in time between two consecutive delay sample determining the RTT of the connection.

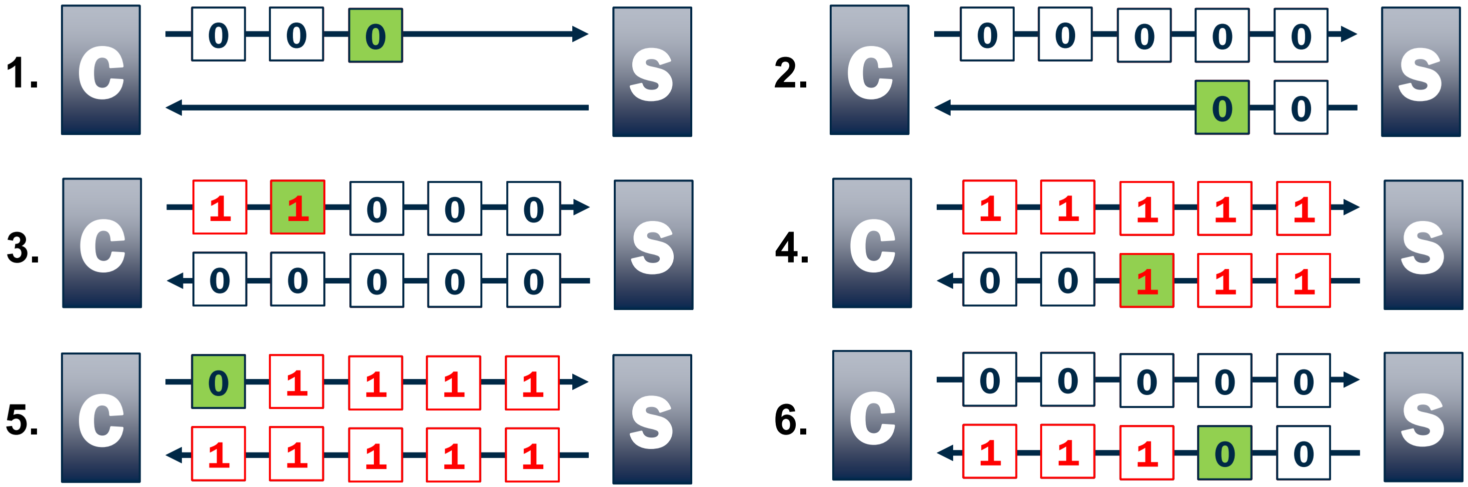

Delay Bit: how the Marked sample works

- **Generation:** when the connection starts, the client set the delay bit of the first packet to 1
- **Reflection:** both endpoints reflect an incoming delay sample to the first outgoing packet
  - If reflection is delayed for more than 1ms (due to lack of traffic), reflection is aborted
- **Client side control:** if a spin-bit period ends without a delay sample
  - the **recovery process** is triggered:
    - the client waits an **empty period** in which no delay sample is introduced;
    - then, it regenerate the delay sample marking the first packet of the following spin-bit period.

- The empty period is needed to signal to possible observer that there was an issue and a new delay measurement session is starting.

How Delay Bit improves the spin bit mechanism

- **Key Goal: stabilize RTT measurements influenced by packet loss and reordering**
- Packet Loss → already solved by Delay Sample working principles (single sample for period, empty period when it is lost).
- Packet Reordering → has no effects because RTT samples are computed just tracing a single packet, the delay sample.
  - However, the observer must be able to correctly identify periods and the related Delay Sample, as well as empty periods used by client to inform observer that there was a loss or a delay so the sample was discarded. Spurious spin edges generate fake empty spin periods.
  - This can be solved introducing the **waiting interval** into the observer: it is implemented using an interval added after a Spin Bit transition during which any other spin transition is rejected.
- Traffic holes → delayed delay samples are not reflected by the endpoint
  - This trigger the recovery process

Testing platform

- Protocol implementation used: QuicGo
- Network topology (Mininet):

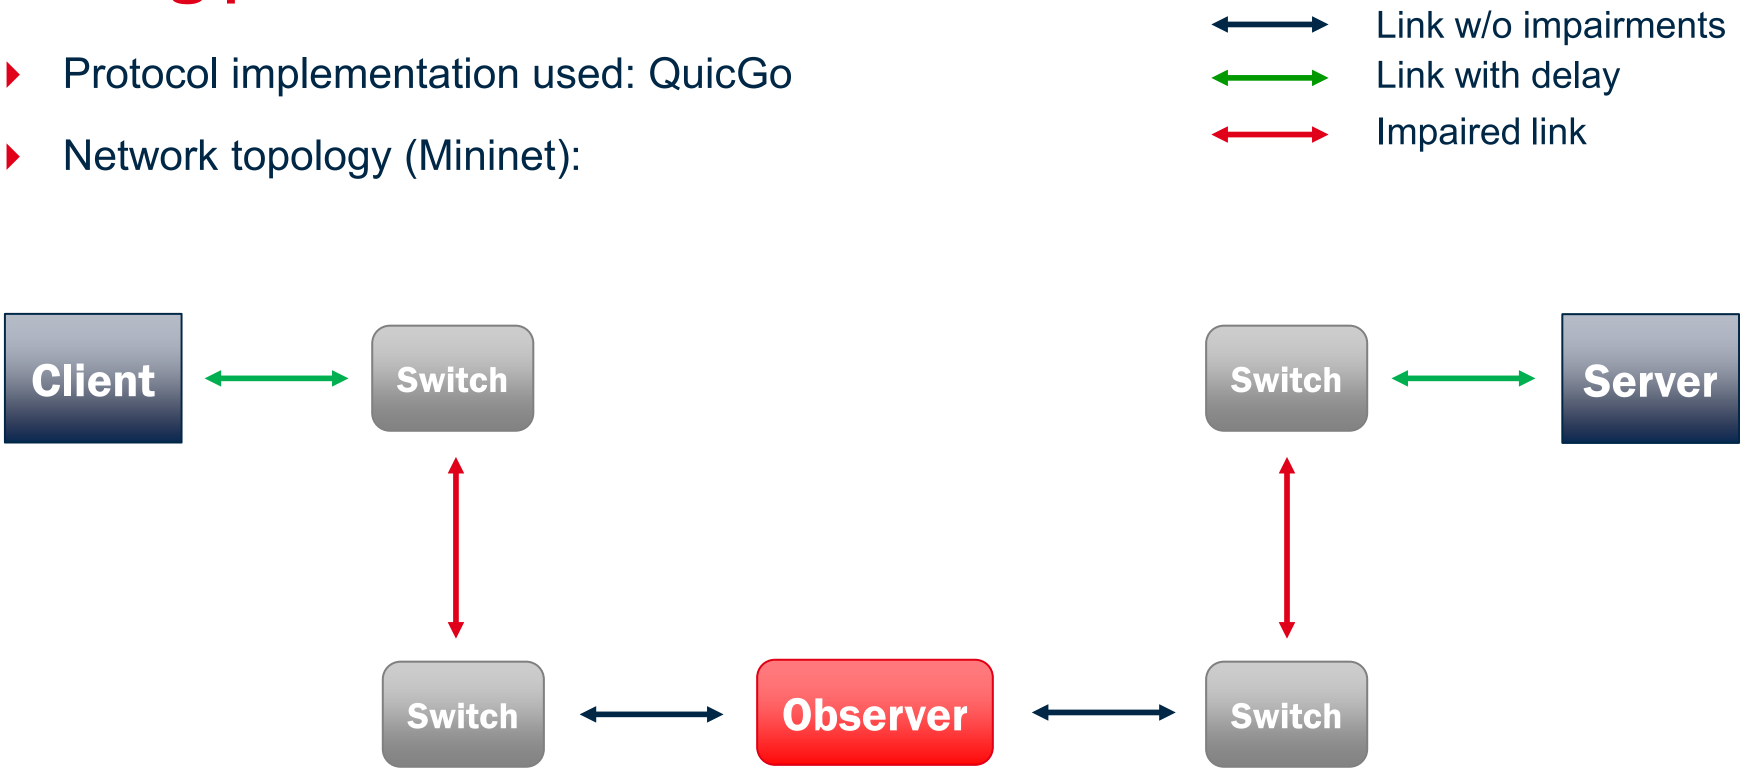

Delay Bit – Test results

- **Delay 40ms**, no loss, no reordering, 200MB stream

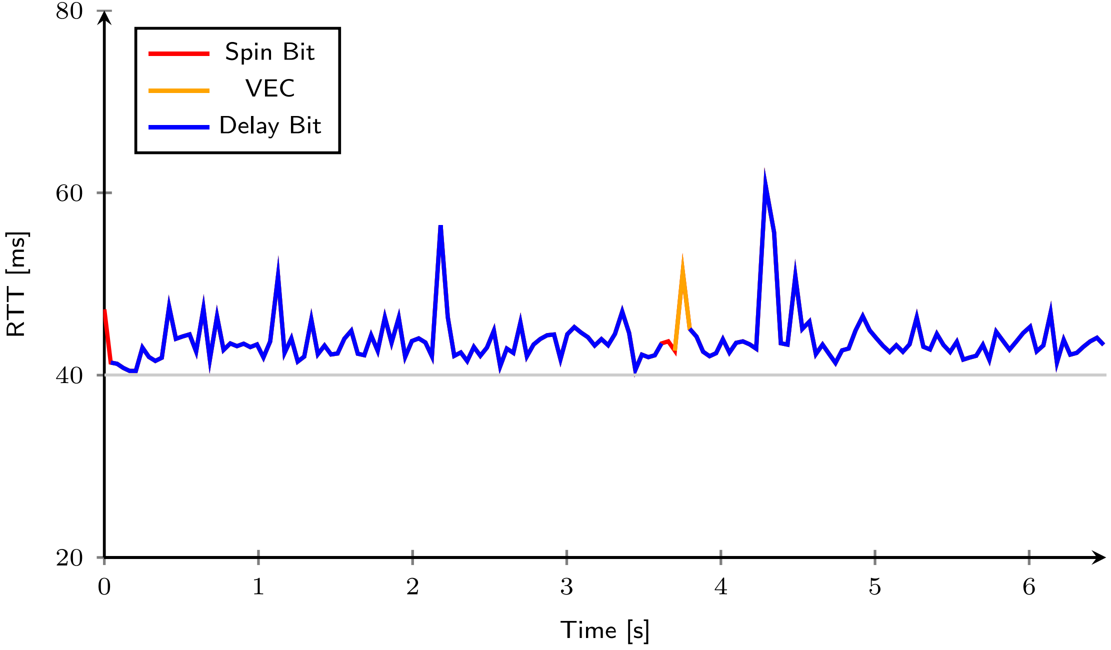

Delay Bit – Test results

- **Delay 40ms**, **random loss**, no reordering, 200MB stream

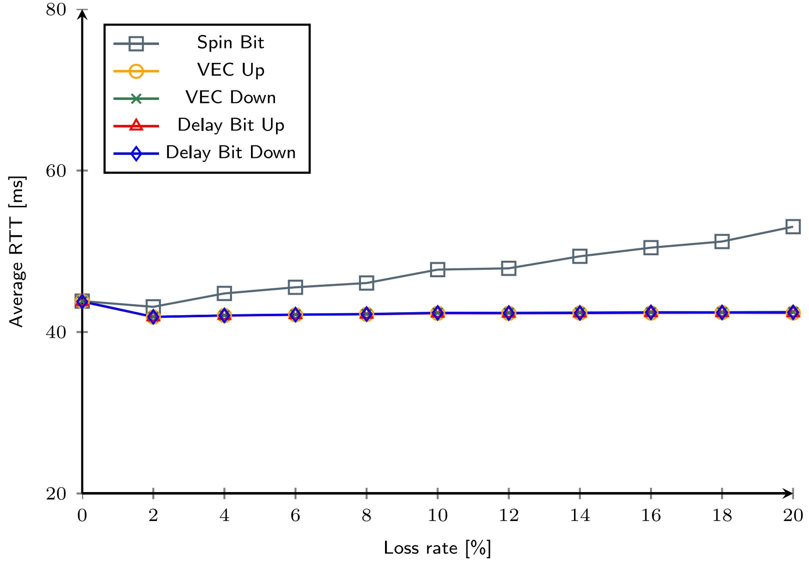

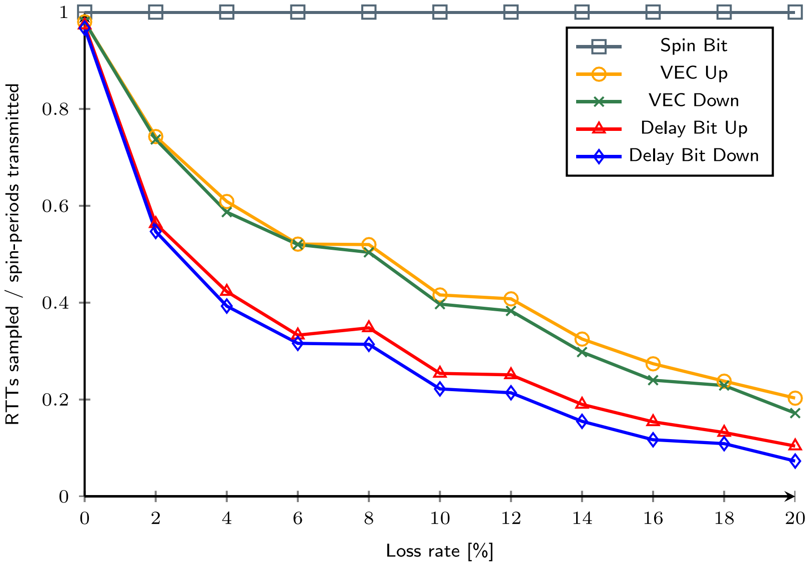

Delay Bit – Test results

- **Delay 40ms**, no loss, **reordering (1ms)**, 200MB stream

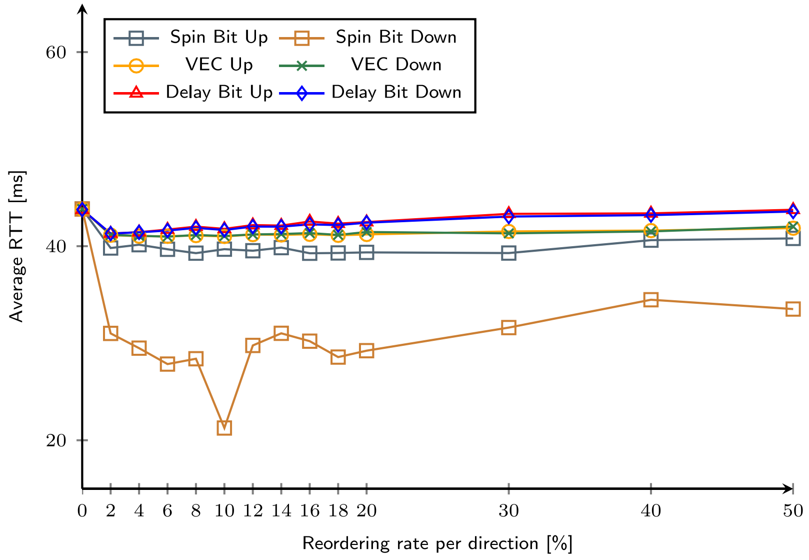

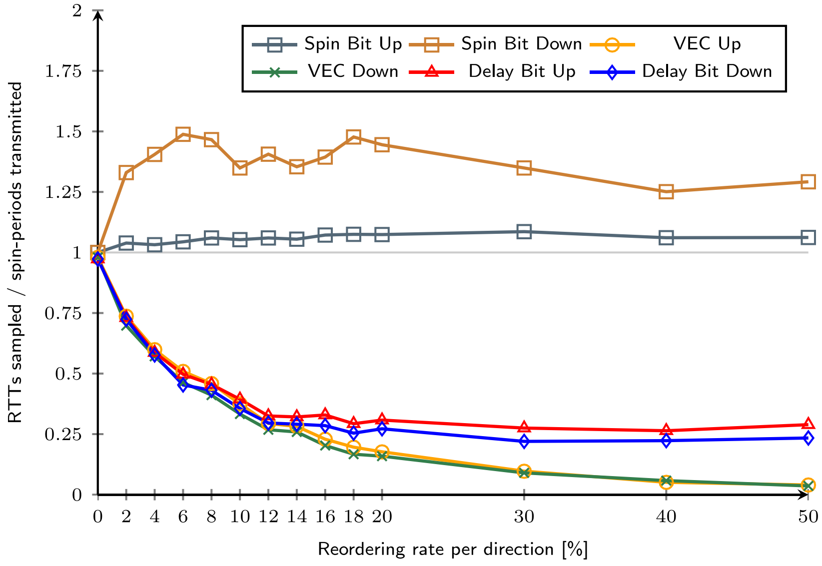

Conclusions

VEC: strengths and weaknesses

- \+ Produces one more valid periods for each edge loss (it's quicker on restart).
- \+ Observer implementation simpler than Delay Sample (no timer).
- \- Requires three bits (the entire amount made available for experimentations).
- \- Decreases its performance in the presence of packet reordering (discarded measurements).

Delay Bit: strengths and weaknesses

- \+ Requires only two bits, leaving the third one available for Packet Loss measurement.
- \+ Produces more valid periods in case of packet reordering (it does not discard periods and produces correct measurements).
- \- Produces less valid periods in case of losses (slower on restart when a delay sample is lost or delayed: an empty period is left).
- \- Observer implementation needs a timer, the waiting period, to skip false periods in case of packet reordering (the waiting period duration is a tradeoff because it is also the minimum measurable RTT).

IETF Drafts

https://tools.ietf.org/html/draft-cfb-ippm-spinbit-new-measurements-01
https://tools.ietf.org/html/draft-trammell-ippm-spin-00
https://tools.ietf.org/html/draft-ietf-quic-spin-exp-01

Thank you