Packet Loss Signaling for Encrypted Protocols

draft-ferrieuxhamchaoui-tsvwg-lossbits

Alexandre Ferrieux – Orange Labs
Isabelle Hamchaoui – Orange Labs
Igor Lubashev – Akamai

IETF 105

Motivation: Loss Detection/Localization Matters

Networks can look like dumb pipes,
<u>only</u> if a plumber can find leaks and patch them quickly

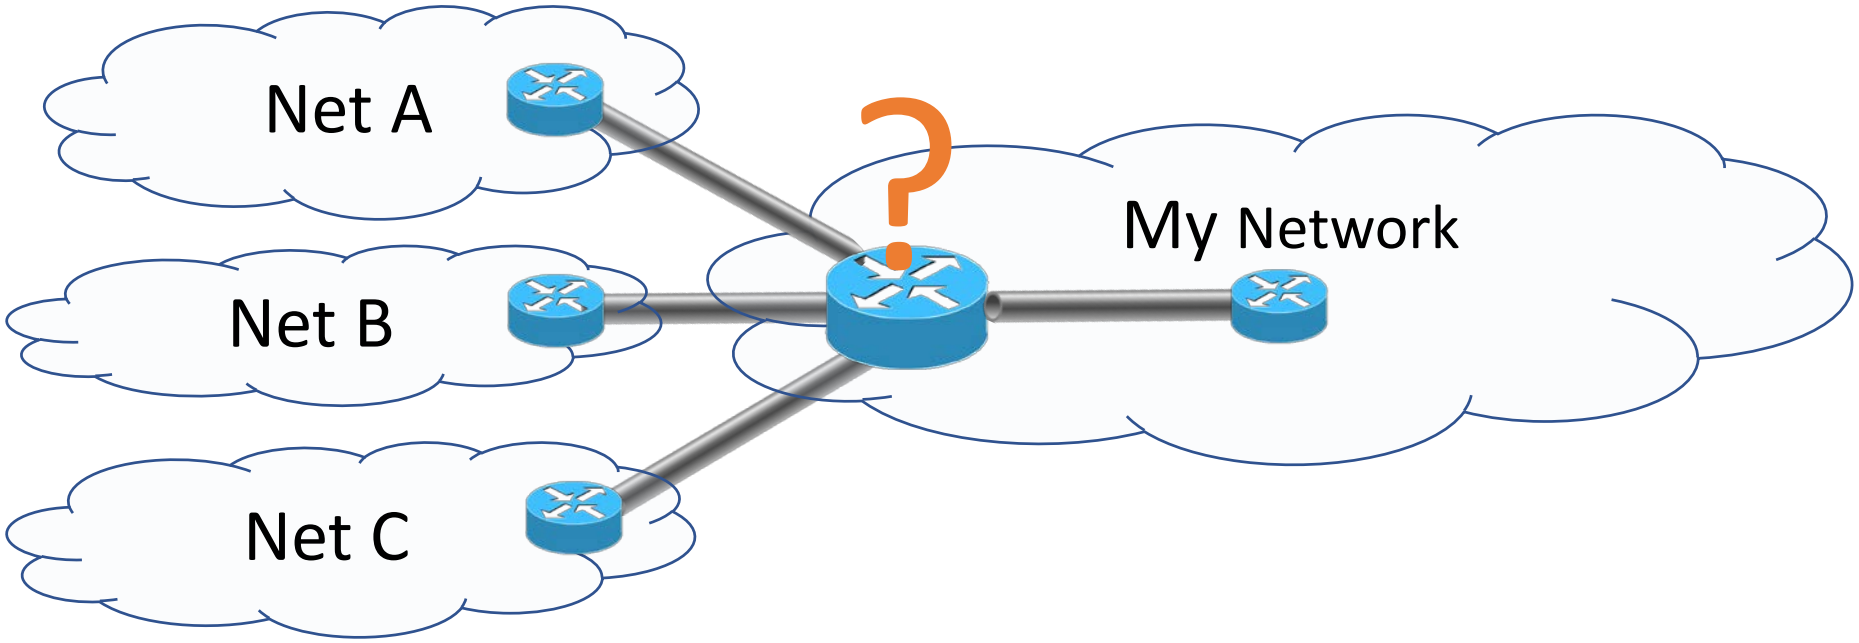

Motivation: Loss Detection/Localization Matters

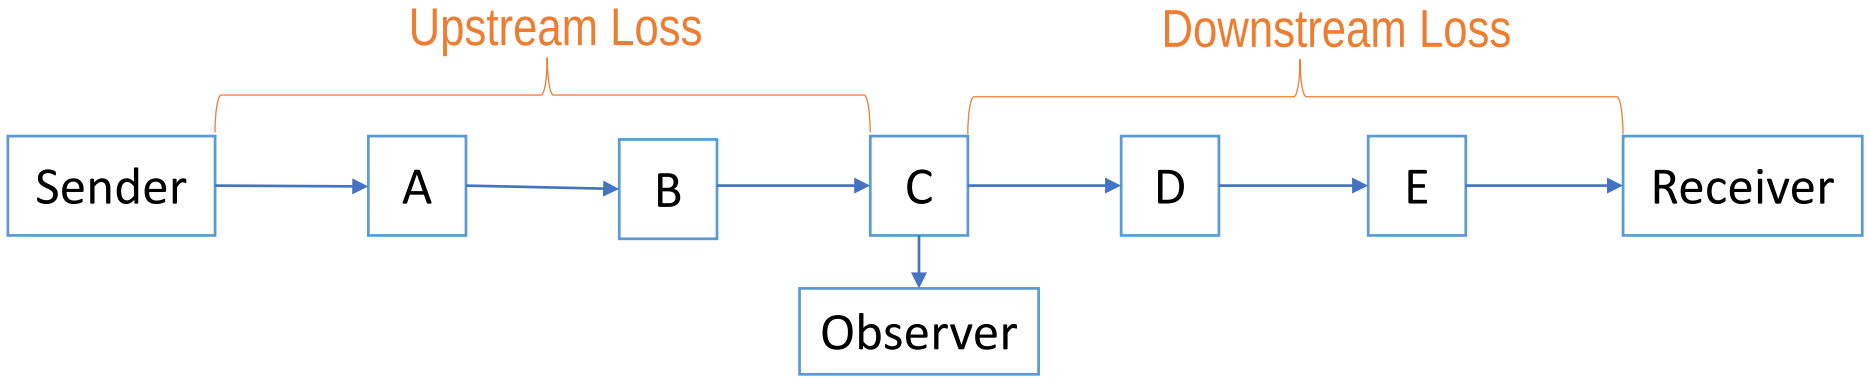

- TCP: observe seq# (and ack#/sack#s, if path is symmetric)
- Transport with encrypted headers: ☹
  - QUIC has a “latency Spin bit”, so you may get an RTT estimate but *not* loss
- *“Just observe similar TCP flows”* is not a good answer

I am not a Network Operator. Why Should I Care?

- If you publish content or services, you derive some benefit from those sites being available and fast.

- If you are a CDN, your customers pay you to ensure their sites are available and fast and to take care of “those Internet issues”.

Proposal: Two “Loss bits”

- **Q**: The “sQuare signal” bit is toggled every N outgoing packets (akin to *color* in RFC 8321)

- **L**: The “Loss event” bit is 1 when “Unreported Loss Counter” (ULC) > 0
  - ULC is incremented for each packet deemed lost by the protocol
  - ULC is decremented for each packet sent with L=1

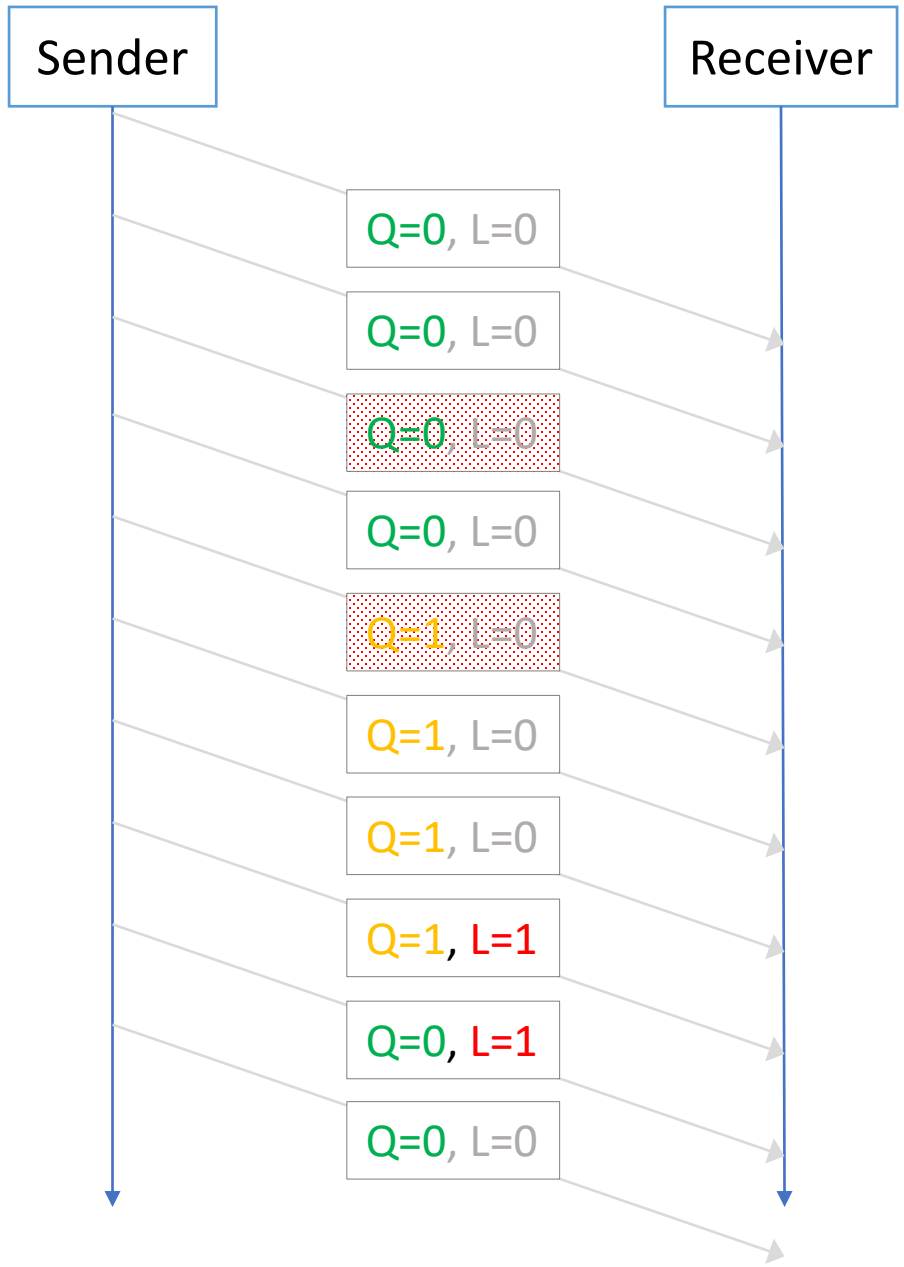

Loss Calculation

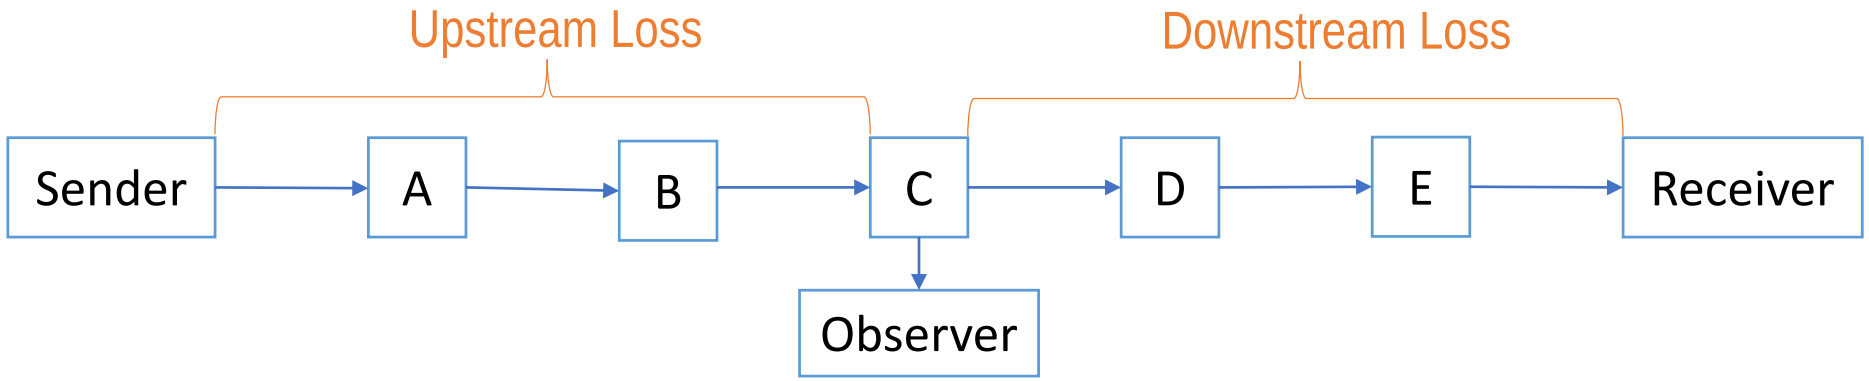

- End-to-End loss (e)

$$e = \text{fraction of packets with L=1}$$

- Upstream loss (u)

$$u = 1 - \frac{\text{average \# of observed packets in a block (same Q)}}{\text{size of the block}}$$

- Downstream loss (d)

$$(1-u)(1-d) = 1-e \qquad d = \frac{e-u}{1-u} \approx e-u$$

Which Protocol Header?

- This draft requires answers to:

  Question 1: “Do we need loss detection by non-endpoints?”

  Question 2: “If we do, are Q & L bits fit for the purpose?”

- If “Yes” to both of the above, we can find a home for the bits in a subsequent draft (possibly in a different WG):
  - IPv4/IPv6 header?
  - IPv4 options / IPv6 HBH option?
  - UDP trailer?
  - QUIC header?

Experimental Data – Akamai to Orange (4 countries)

- Q&L bits are in $ip.ttl \gg 6$ (and $ip6.hoplimit \gg 6$)
- A lot more data and discussion in maprg tomorrow at 10am

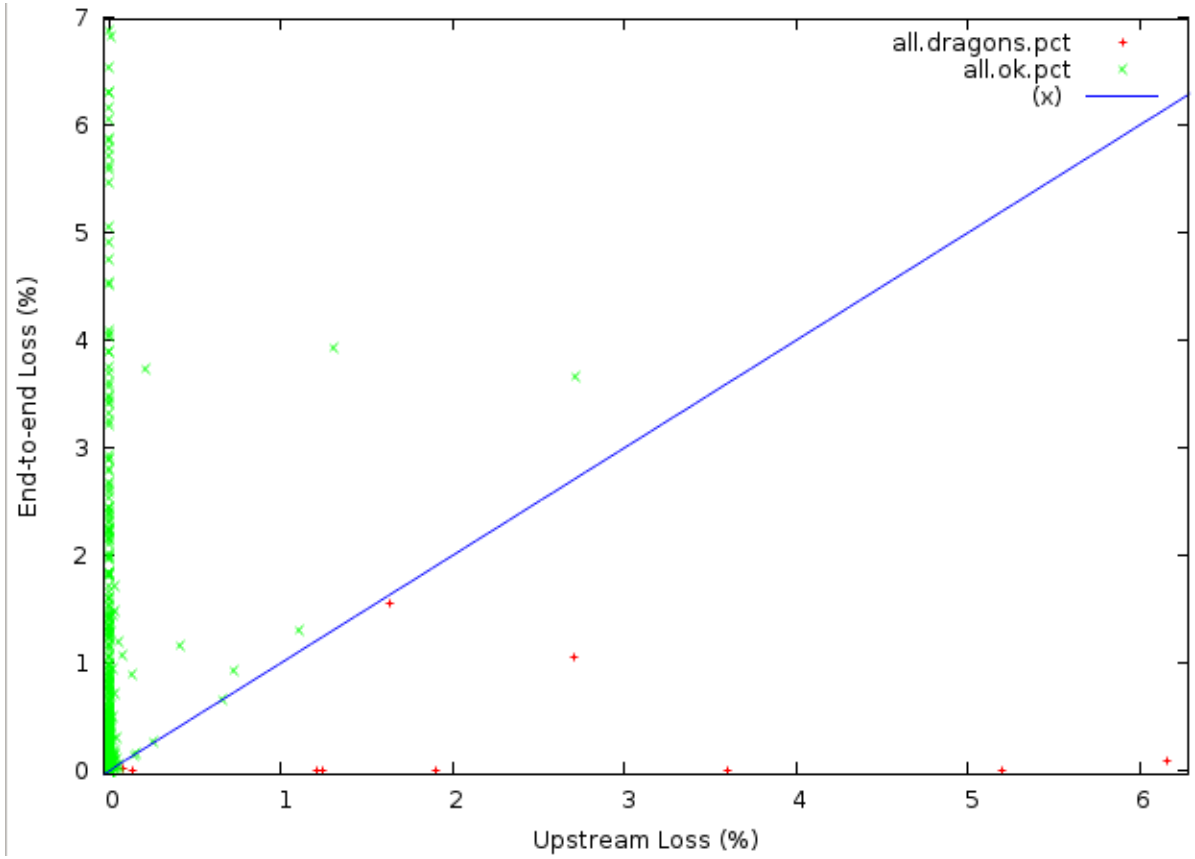

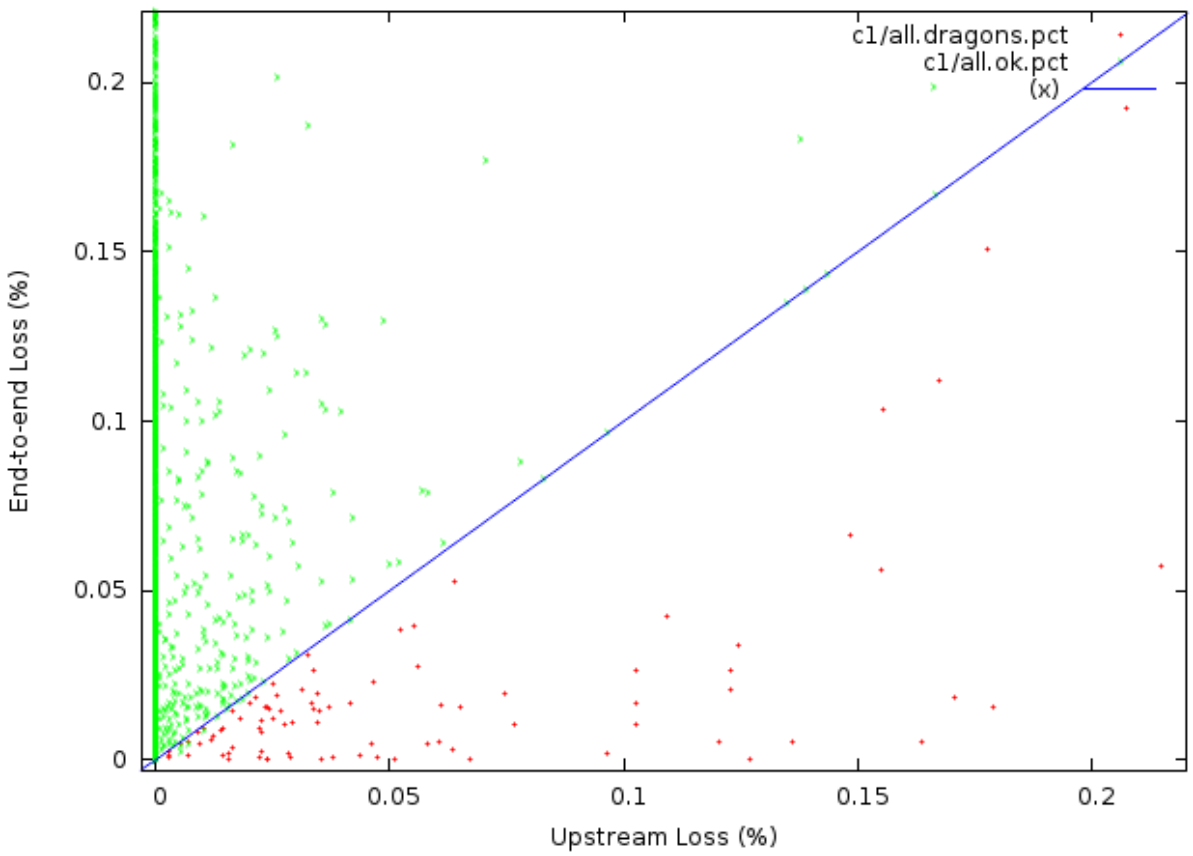

Privacy and Ossification

- Protecting Privacy
  - Explicit signal means less information leakage (RFC 8558)
  - Separate counters for separate flows, subflows, paths, QUIC connection IDs, … to prevent loss signals used to link multiple connections to the same device

- Ossification Resistance
  - Loss signals are not integral protocol bits, so they can be greased, if desired
  - QUIC latency spin bit is an example:
    - can mandate random-looking values for Q&L bits if unused
    - can mandate to not using for a certain portion of connections

Getting in Touch

- Mailing List: ietf-loss-bits@googlegroups.com
- Data Discussion on Friday at 10am (mapgr)
- draft-ferrieuxhamchaoui-tsvwg-lossbits