

Spurious Retransmission Detection (SRD) with the TCP Echo Options

`draft-zimmermann-tcpm-spurious-rxmit`

Richard Scheffenegger <`rs@netapp.com`>

Alexander Zimmermann <`alexander.zimmermann@netapp.com`>

Problem Statement

■ Eifel detection

- Uses TCP Timestamp options [RFC 7323] to detect spurious retransmissions
- Limited applicability due to TSecr semantics, and TSval granularity
- No detection of reordering during loss recovery

■ Idea: Make every segment – including all retransmissions – uniquely identifiable (to the sender)

- Allows all functionality of Eifel, even during corner cases
- Enables new capabilities (lost retransmission detection)

Spurious Retransmission Detection (SRD)

■ Mechanism



- Use TCP Echo Option to send a (small) counter in each segment to keep MSS equal for retransmissions
- Increase counter when sending a new round of retransmissions e.g. (re-)entering loss recovery
- Check counter in received ACK
 - Equal to current value → valid retransmission
 - Else spurious retransmissions

■ Property

- Semantics of TCP Echo allows to determine the exact ordering of transmissions, even in case of reordering



Example (Semantics)

- RFC7323 TSecr reflects TS of last in-sequence segment



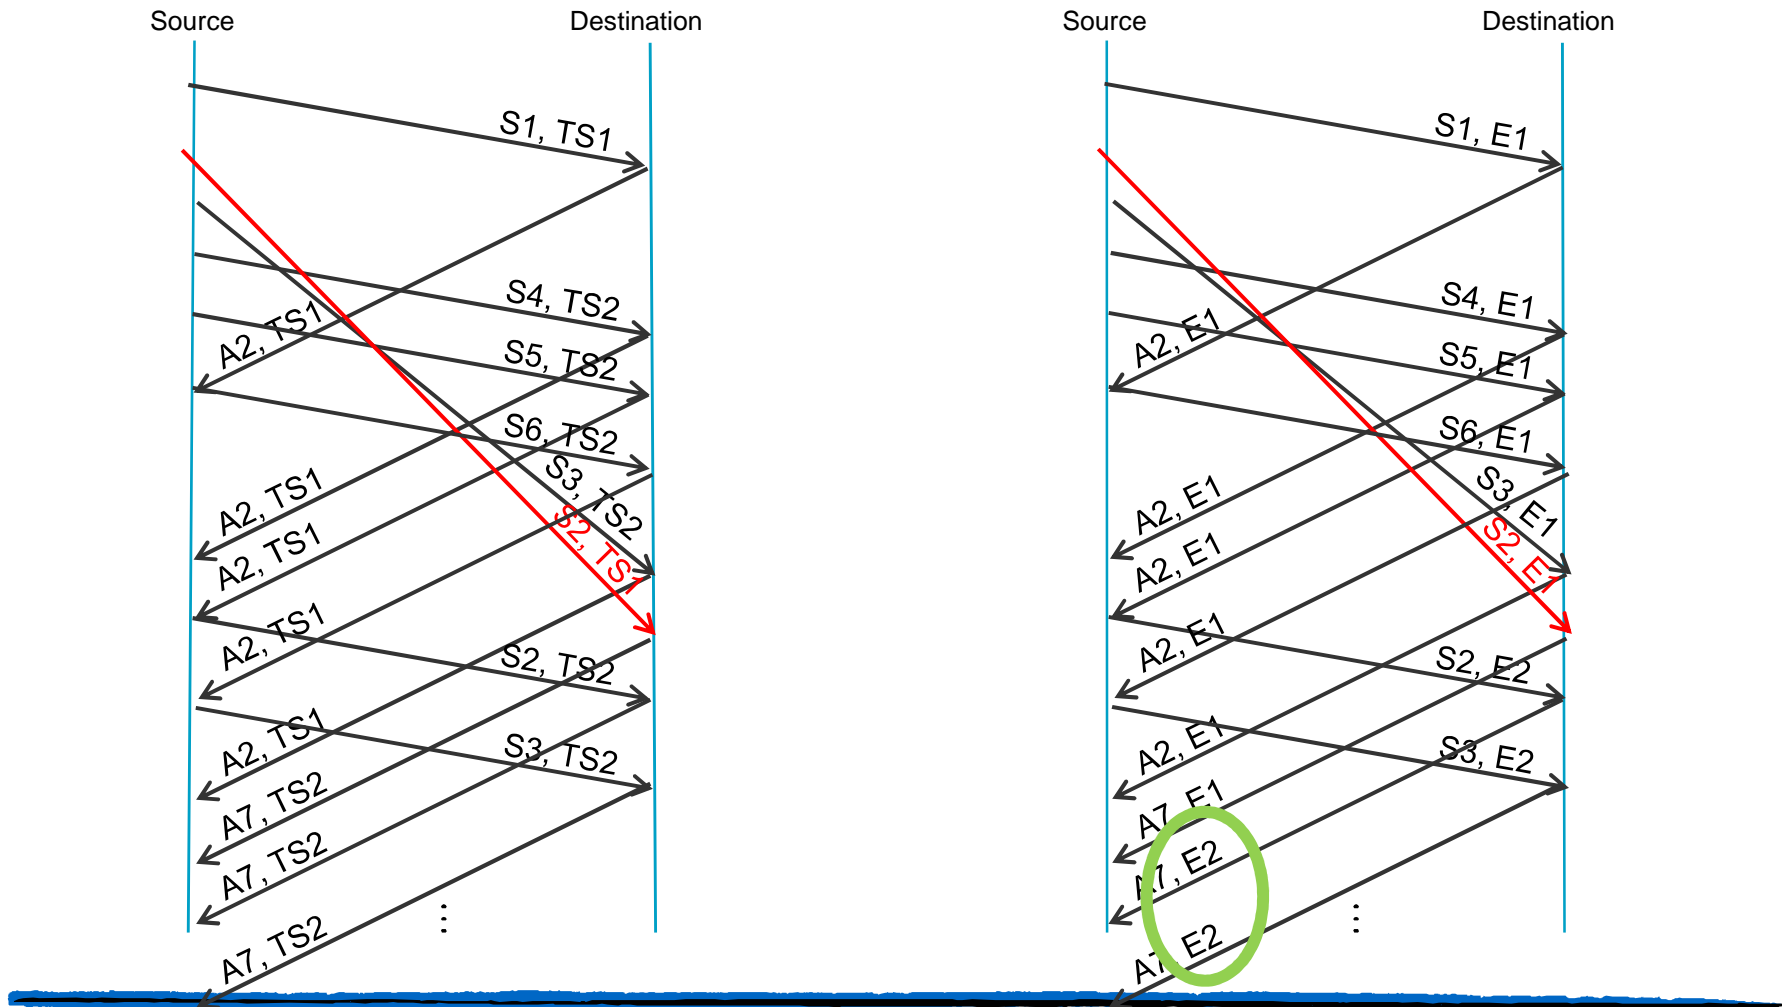
Example (Eifel vs. SRD)

- Granularity of TS often too coarse




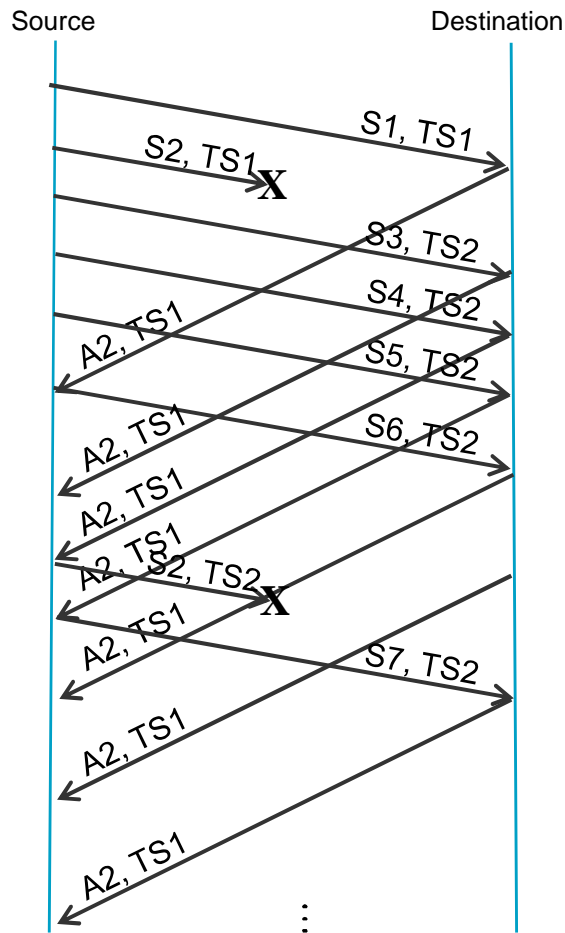
Example (Eifel vs. SRD)

- Eifel only works on first retransmitted segment



Example (Eifel vs. SRD)

- Allows lost retransmission detection



Moving forward...

- **Less overhead than RFC7323 Timestamps**
- **Solves the retransmission ambiguity problem completely**
 - More Complex scenarios involving Fwd Loss / Fwd Reordering / ACK Loss / ACK Reordering
 - Enables Lost Retransmission Detection (LRD) while strictly adhering to packet conservation principles
 - QUIC has similar “control sequence number”
- **Next steps**
 - Received initial feedback (clarifications)
 - Eventually asking for adoption