Robust Channels

Handling Unreliable Networks in the Record Layers of QUIC and DTLS 1.3

Marc Fischlin, **Felix Günther**, Christian Janson

TECHNISCHE UNIVERSITÄT DARMSTADT

ETH zürich

DFG Deutsche Forschungsgemeinschaft
German Research Foundation

QUIC/DTLS 1.3 within the Network Stack

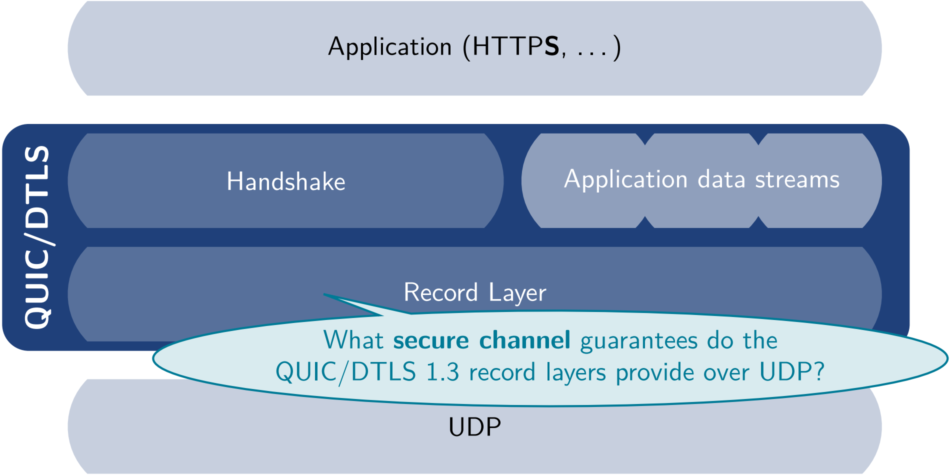

What **secure channel** guarantees do the QUIC/DTLS 1.3 record layers provide over UDP?

Recap: Secure Channels over TCP

. . . think: TLS

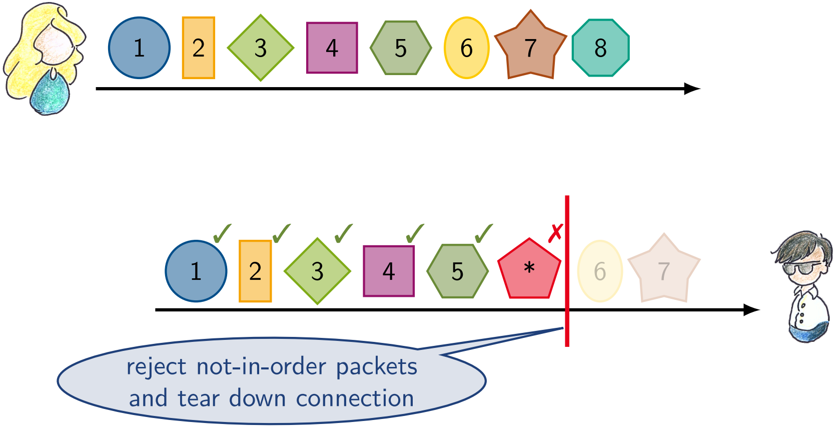

reject not-in-order packets and tear down connection

drawings by *Giorgia Azzurra Marson*

Handling Unreliable Transport

QUIC, DTLS, . . . over UDP

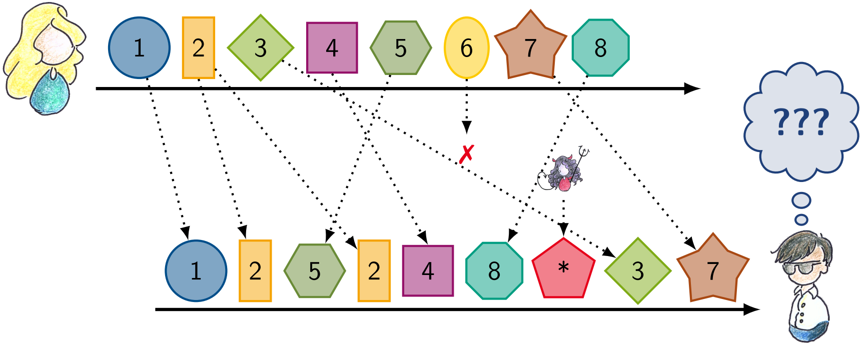

Handling Unreliable Transport

Many choices. . .

- **Replays / Duplicates**
  - prevent them?
  - check how far back?

	QUIC	DTLS 1.3
prevent them?	MUST prevent	optional
check how far back?	e.g., anti-replay window (IPsec)	

- **Reordering**
  - permitted?
  - by how far max.?

	QUIC	DTLS 1.3
permitted?	. . . well, yes—it's UDP . . .	
by how far max.?	dynamic 1–4B window	dynamic 1–2B window

- **Adversarial interaction**
  - Integrity: reject non-genuine packets

	QUIC	DTLS 1.3
Integrity: reject non-genuine packets	rely on AEAD	

- But how do you (formally) guarantee that replayed / reordered / adversarial packets don't affect others?

new notion: **Robustness**

Generalizing Channel Correctness

. . . beyond prior hierarchies [BKN02,KPB03,Boy+16,RZ18]

- parameterize what packet (ciphertext) reordering a channel **supports**
- predicate $\mathsf{supp}(C_S, C_R, c) = ✓ \,/\, ✗$
  - C_S: sequence of sent ciphertexts
  - C_R: sequence of *supported* ciphertexts received prior
  - c: next ciphertext to receive
- correctness (only) requires genuine, supported ctxts be correctly decrypted

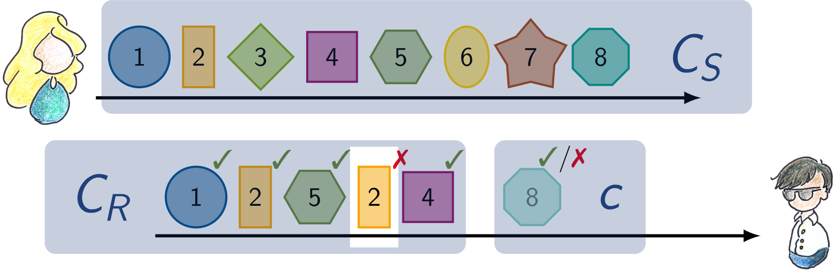

Defining Robustness (ROB)

“malicious packets cannot disturb expected channel behavior”

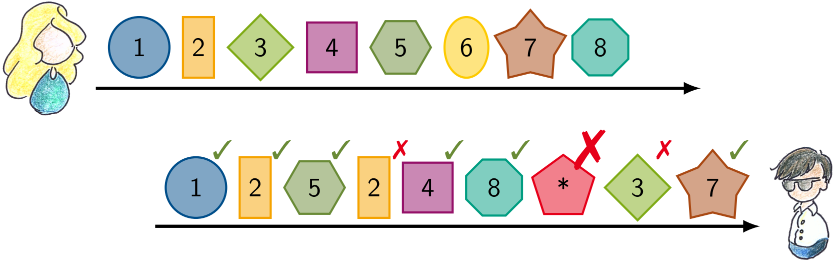

Defining Robustness (ROB)

Idea: Compare with the supported, correct sub-trace

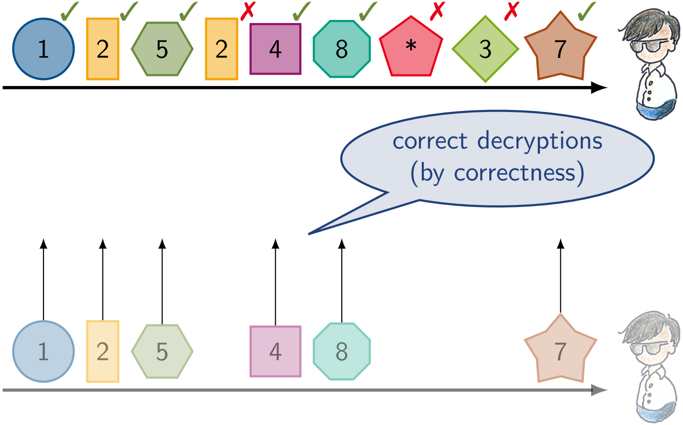

Defining Robustness (ROB)

Idea: Compare with the supported, correct sub-trace

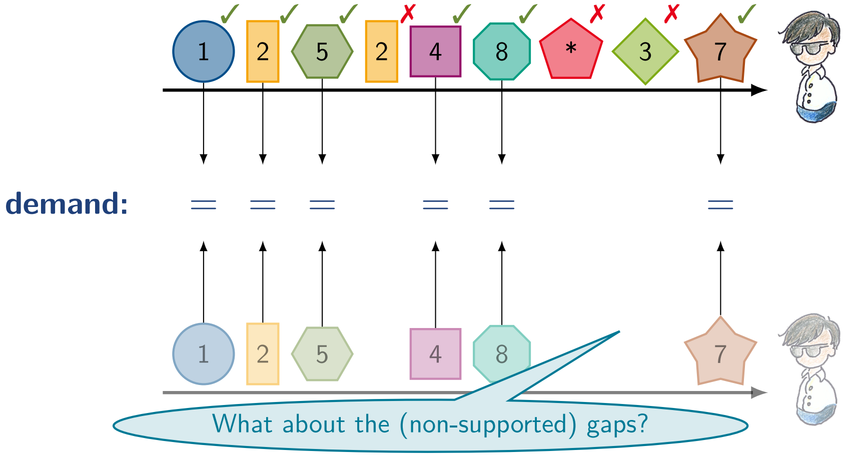

Robust Integrity (ROB-INT)

- join **robustness** and **integrity** for desired property over unreliable transport

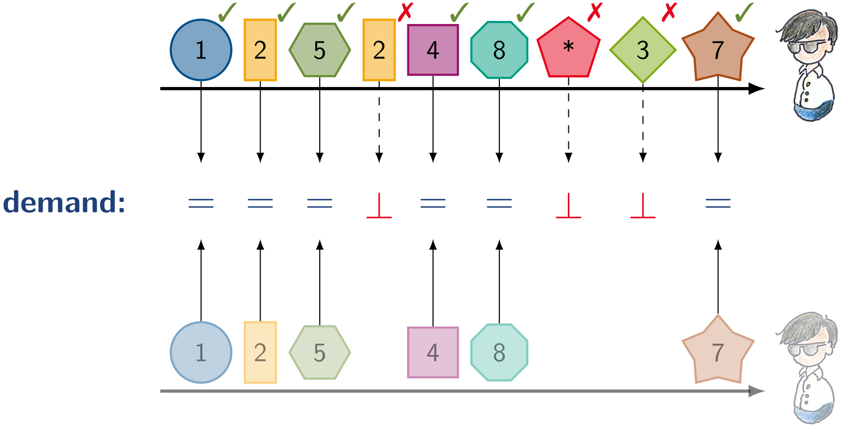

A Robust Hierarchy

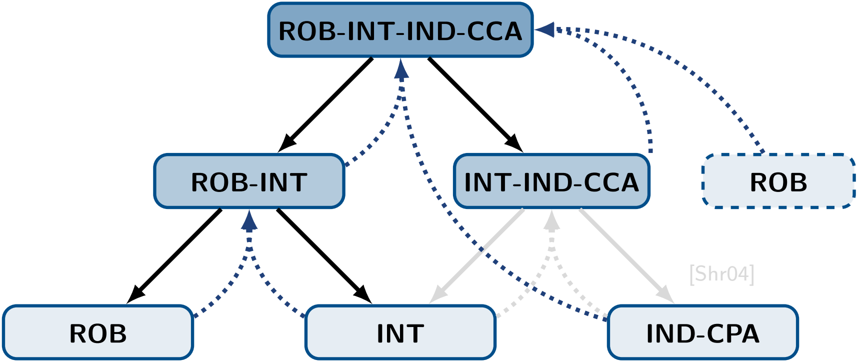

QUIC Channel

Correctness for Dynamic Sliding Windows

ETH zürich

- header (w/ partial packet no. pn_e) + AEAD ciphertext
- pn_e defines $|pn_e|$-**bit dynamic sliding window**
- check for **replays in w_r-sized window**

.. | $pn_e = 000$ | ..

AEAD
ctxt

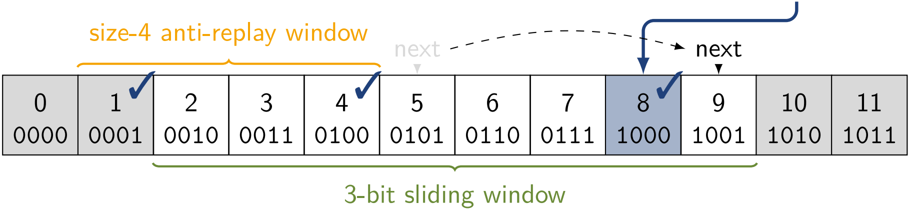

0	1	2	3	4	5	6	7	8	9	10	11
0000	0001	0010	0011	0100	0101	0110	0111	1000	1001	1010	1011

$$\mathsf{supp}_{dw\text{-}r[w_r]}(C_S, C_R, c) :=$$

supported if in sliding window (dynamic for c) **and** replay window

$$\Big[c \in C_S \land c \notin C_R \land \mathsf{index}(c, C_S) \in [\mathsf{n} - \min(w_b^c, w_r + 1), \mathsf{n} + w_f^c] \Big]$$

(simplified)

QUIC Channel

Robust Confidentiality and Integrity (ROB-INT-IND-CCA)

- use hierarchy: **IND-CPA + ROB-INT = ROB-INT-IND-CCA**

$$\mathrm{Adv}^{\text{ROB-INT-IND-CCA}}_{\text{QUIC}} \leq \mathrm{Adv}^{\text{priv}}_{\text{AEAD}} + q_R \cdot \mathrm{Adv}^{\text{auth}}_{\text{AEAD}}$$

  - important: can make **multiple forgery attempts**
  - factor q_R (#received ciphertexts) loss in security reduction

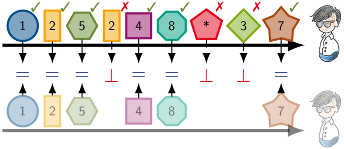

Responses

- IETF WGs updated QUIC / DTLS 1.3 drafts to mandate **concrete forgery limits** (beyond confidentiality limits [LP17])

> The integrity protections ... depend on limiting the number of attempts to forge packets. ... QUIC ignores any packet that cannot be authenticated, allowing multiple forgery attempts.

- **Usage Limits on AEAD Algorithms** `draft-irtf-cfrg-aead-limits`
  - new CFRG document draft (w/ Chris Wood, Martin Thomson)
  - aims to provide user guidance on AEAD usage limits
  - confidentiality/integrity, single-/multi-key, AES-GCM/AES-CCM/ChaCha20Poly1305

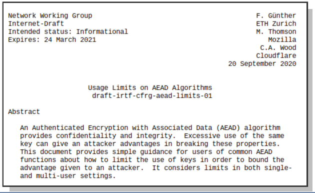

```
Network Working Group                                         F. Günther
Internet-Draft                                                ETH Zurich
Intended status: Informational                                M. Thomson
Expires: 24 March 2021                                           Mozilla
                                                                C.A. Wood
                                                              Cloudflare
                                                       20 September 2020


                      Usage Limits on AEAD Algorithms
                       draft-irtf-cfrg-aead-limits-01

Abstract

   An Authenticated Encryption with Associated Data (AEAD) algorithm
   provides confidentiality and integrity.  Excessive use of the same
   key can give an attacker advantages in breaking these properties.
   This document provides simple guidance for users of common AEAD
   functions about how to limit the use of keys in order to bound the
   advantage given to an attacker.  It considers limits in both single-
   and multi-user settings.
```

Summary

- We introduce **robustness** as first-class security property

  “malicious packets cannot disturb expected channel behavior”

- We analyze **QUIC and DTLS 1.3**
  - capturing dynamic sliding window & replay-checking
  - confirm both achieve intended robust confidentiality and integrity
  - . . . but q_R loss has to be taken into account

- Led to **updated QUIC and DTLS 1.3 drafts**, mandating forgery limits

full version @ IACR ePrint: `https://ia.cr/2020/718`

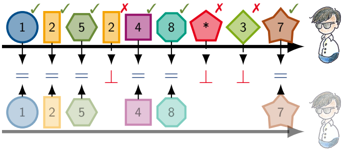

Thank You!

mail@**felixguenther.info**

References I

[BKN02] M. Bellare, T. Kohno, and C. Namprempre. "Authenticated Encryption in SSH: Provably Fixing The SSH Binary Packet Protocol". In: *ACM CCS 2002*. Ed. by V. Atluri. ACM Press, Nov. 2002, pp. 1–11.

[Boy+16] C. Boyd, B. Hale, S. F. Mjølsnes, and D. Stebila. "From Stateless to Stateful: Generic Authentication and Authenticated Encryption Constructions with Application to TLS". In: *CT-RSA 2016*. Ed. by K. Sako. Vol. 9610. LNCS. Springer, Heidelberg, 2016, pp. 55–71.

[KPB03] T. Kohno, A. Palacio, and J. Black. *Building Secure Cryptographic Transforms, or How to Encrypt and MAC.* Cryptology ePrint Archive, Report 2003/177. http://eprint.iacr.org/2003/177. 2003.

[LP17] A. Luykx and K. G. Paterson. *Limits on Authenticated Encryption Use in TLS.* http://www.isg.rhul.ac.uk/~kp/TLS-AEbounds.pdf. Aug. 2017.

[RZ18] P. Rogaway and Y. Zhang. "Simplifying Game-Based Definitions - Indistinguishability up to Correctness and Its Application to Stateful AE". In: *CRYPTO 2018, Part II*. Ed. by H. Shacham and A. Boldyreva. Vol. 10992. LNCS. Springer, Heidelberg, Aug. 2018, pp. 3–32.

[Shr04] T. Shrimpton. *A Characterization of Authenticated-Encryption as a Form of Chosen-Ciphertext Security.* Cryptology ePrint Archive, Report 2004/272. http://eprint.iacr.org/2004/272. 2004.