Network Visibility in the Era of Encrypted Client Hello

Richard Plny, Jan Luxemburk, Jaroslav Pesek
CESNET, Czechia

TLS 1.3 Encrypted Client Hello (ECH) Extension

- Encrypts information from TLS Client Hello (CH)
  - **Server Name Indication** (SNI) !
- A new extension (ECH) with encrypted data is appended
  - Contains real SNI, for example *discord.com*
- Original fields in TLS CH contain non-sensitive data
  - Cloudflare puts *cloudflare-ech.com* into original SNI

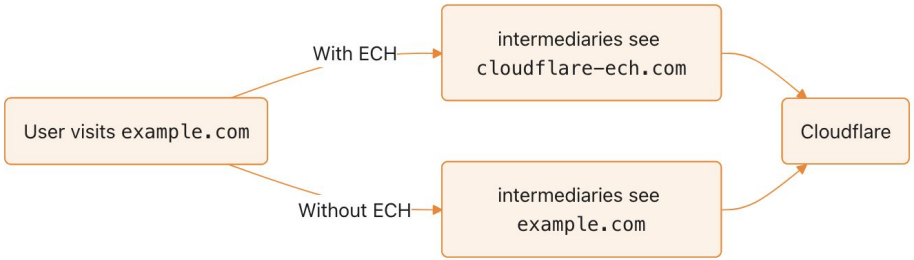

Also in QUIC & HTTP/3

- ECH is in TLS v1.3 and TCP
- But also in QUIC
  - QUIC has TLS integrated
- ECH config is distributed by:
  - DNS HTTP records
  - Directly sent by the server (HelloRetryRequest)
  - Completely outside the box (can be hardcoded)

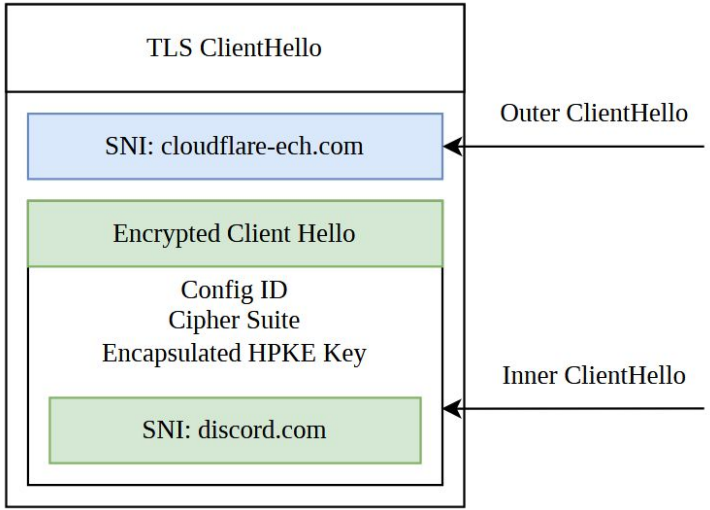

TTL:

30 minutes

DATA:

HTTPS

```
1 . ipv4hint="213.108.108.101" ech="AMD+DQA8GgAgACDs5uweq5/qNDSP/racm8F3A+RHJdhfPQ1rVIW+ogK/
LgAEAAEAAQANY292ZXIuZGVmby5pZQAA/
g0APAIAIAAg4tHHPbADdk+khqVXNG+4+UIphGl5EcvuLnKB9jPM2z0ABAABAAEADWNvdmVyLmRlZm8uaWUAAP4NADwVACAAIMp+Swv
ZZitB6YL9GjuI1p+CT2N1w/E8WQUpqYujO8M4AAQAAQABAA1jb3Zlci5kZWZvLmllAAA="
ipv6hint="2a00:c6c0:0:116:5::10"
```

How to classify ECH-enabled traffic?

- Can we use non-ECH enabled connections to reveal traffic to the same SNI?
- How different is ECH and non-ECH traffic?

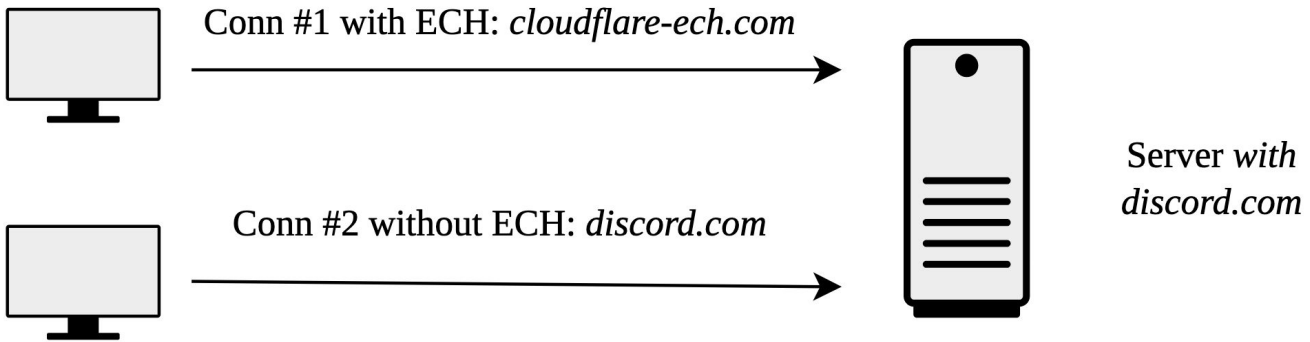

ECH GREASE

- **G**enerate **R**andom **E**xtensions **A**nd **S**ustain **E**xtensibility (GREASE)
- Put ECH extension into CH even when ECH is **not** supported
- Contains dummy data
- When server/middlebox does not support ECH
  - Treated as unknown extension
  - Will not be used
- Bridges the gap between (non-)ECH-enabled traffic

ECH GREASE Support by Web Browsers

- ECH GREASE enabled by default in majority (**~80%**) of browsers
  - Chrome/Chromium, Firefox, Edge, Opera
  - Samsung Internet based on Chromium
- Not supported by Safari
- CH sizes **were already** affected

Client Hello Sizes

- ECH and ECH GREASE affect CH sizes
- CH can also be split into more frames

	Home		Office	
Settings	Chrome	Firefox	Chrome	Firefox
ECH	1703 (186)	1840 (281)	1705 (186)	1840 (281)
GREASE	1763 (250)	1892 (281)	1793 (282)	1892 (281)
NONE	—	1607 (0)	—	1607 (0)

Previous Research - ECH Providers

- [1] Mücke et al. - ***Towards a Complete View of ECH Deployments***
- Currently the largest ECH provider - Cloudflare
  - 2.9 million domains
- Followed by Meta
  - 4.7k
- Some small others

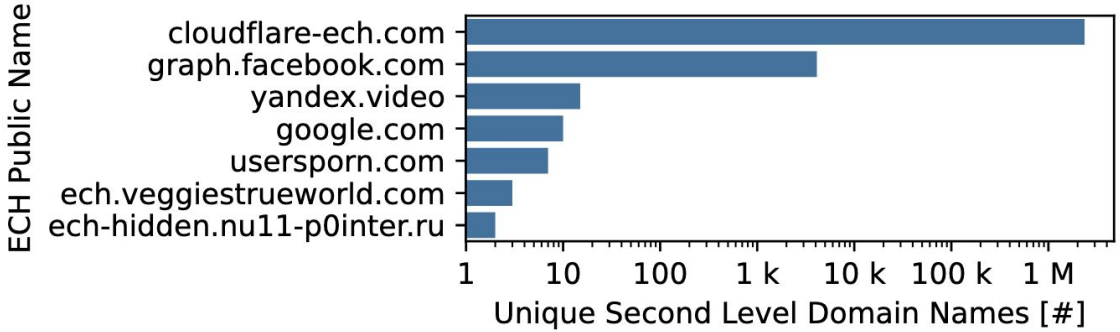

Previous Research - ECH Providers

- [1] Mücke et al. - ***Towards a Complete View of ECH Deployments***
- ECH Discovery
  - Send ECH with random data
  - Server cannot decrypt the ECH
  - Server sends HelloRetryRequest with current ECH config (**ECH supported**)
  - Otherwise ECH is not supported

Figure 2: Simplified TLS handshake with ECH in our measurement setup. We use the HelloRetryRequest from the server to determine ECH-capable deployments.

Idea

- Use non-ECH flows for classification of ECH-enabled flows

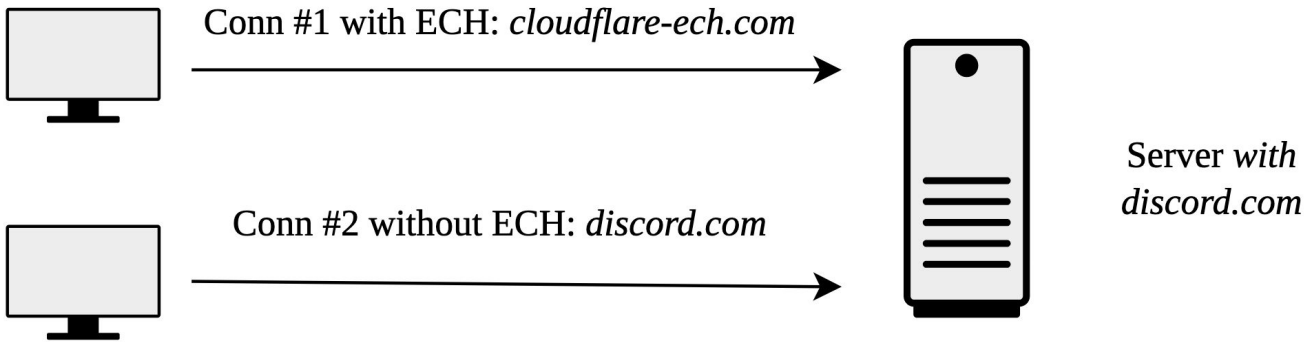

Previous Research - QUIC Domain Recognition

- [2] Luxemburk et al. - ***Universal Embedding Function for Traffic Classification via QUIC Domain Recognition Pretraining: A Transfer Learning Success***
- Network flow → Embedding
  - Embedded into latent space through neural network
  - Use k-NN to assign domains based on similarity
- Before wider ECH deployment

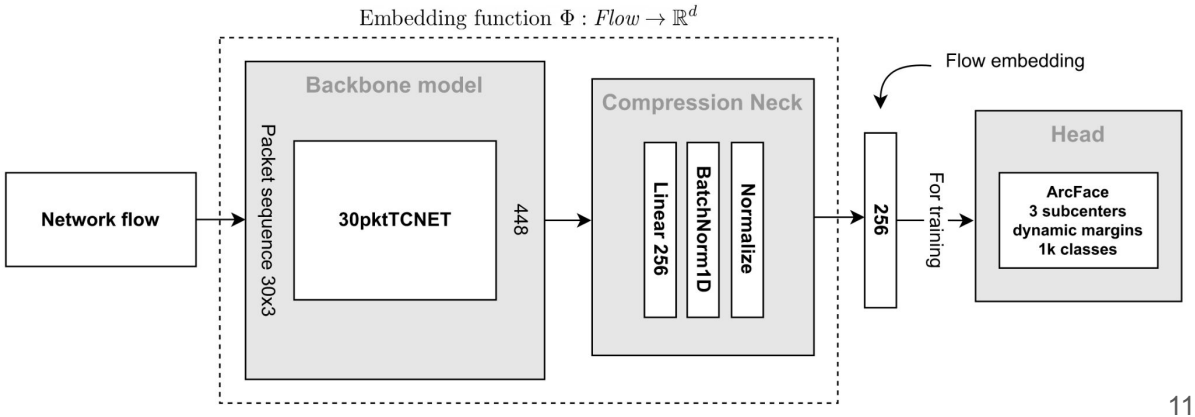

CESNET-QUICEXT-25

- New dataset with QUIC traffic from CESNET backbone lines
- Collected through June 2024 to April 2025
  => contains ECH and ECH GREASE traffic!
- Available on Zenodo
  - https://zenodo.org/records/17249078

Transfer to ECH-enabled traffic

- Thanks to ECH GREASE, CH sizes already affected
- Majority of web browsers use ECH GREASE by default
- Small portion of traffic remained as it was before
- We only use:
  - Flows with destination ASN being **Cloudflare**
  - Flows with domains that **support** ECH

=> We are evaluating the previous method for ECH-enabled QUIC traffic

Network Flow Embeddings

- Network flow - contains information about the first 30 packets
  - Packet sizes, directions, and inter-arrival packet times
    => feature vector with shape 30x3
- 30pktTCNET - neural network
  - Frozen, no fine-tuning

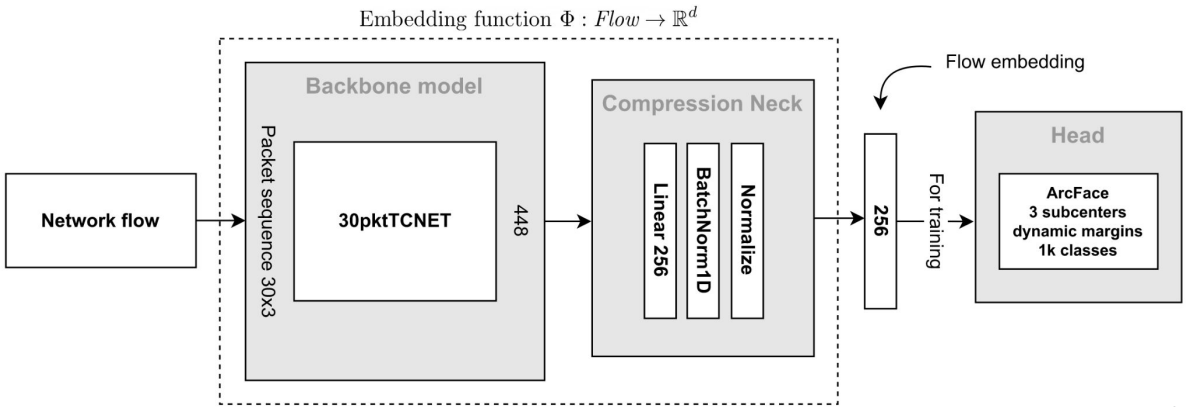

Database & Flow Classification

- Embeddings of non-ECH flows put into database with their domains
- Classification of ECH flow
  - **Flow → Embedding → DB lookup → Domain**
  - Domain can be assigned in multiple ways (next slide)
- First step: evaluation of the method on new dataset on non-ECH flows
  - But with ECH GREASE data
  - Our goal: transfer it for ECH flows as well

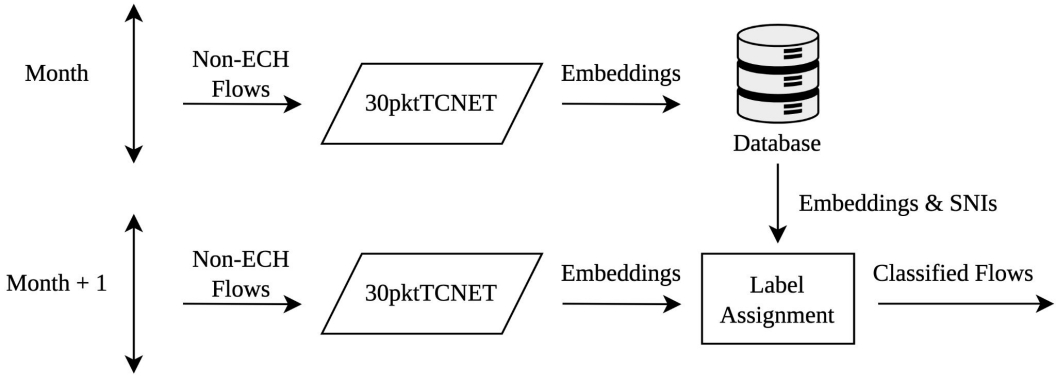

Multiple Methods of Domain Assignment

- TOP1
  - Assign the domain of the closest embedding in the db
- MAJ-*K*
  - Assign a majority domain of the *K* closest embeddings
- SMART-MAJ
  - Select all embeddings until certain distance threshold
  - Assign a majority domain from such embeddings

Results on new data with ECH (GREASE)

- Only ECH-supported TOP50 domains from Cloudflare ASN

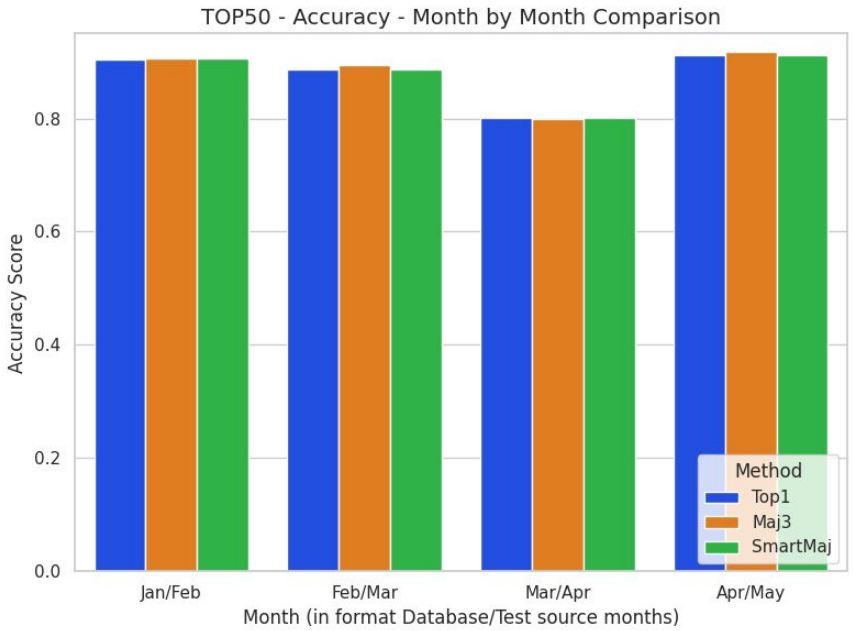

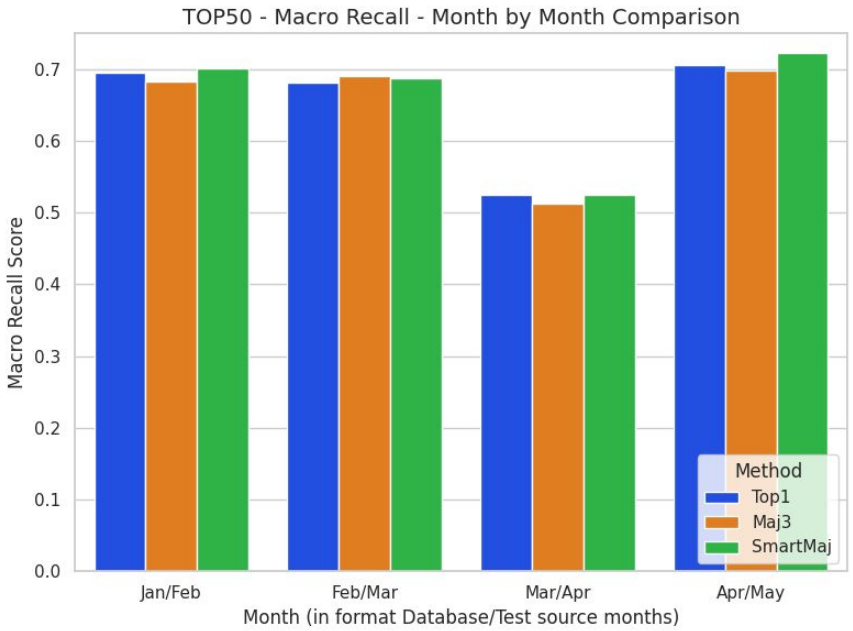

Results on new data with ECH (GREASE)

- Only ECH-supported TOP100 domains from Cloudflare ASN

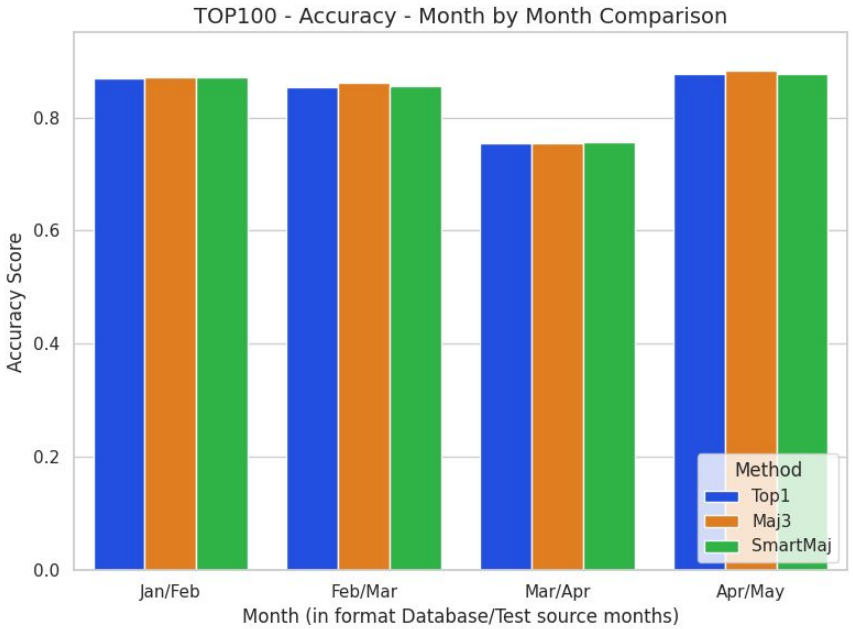

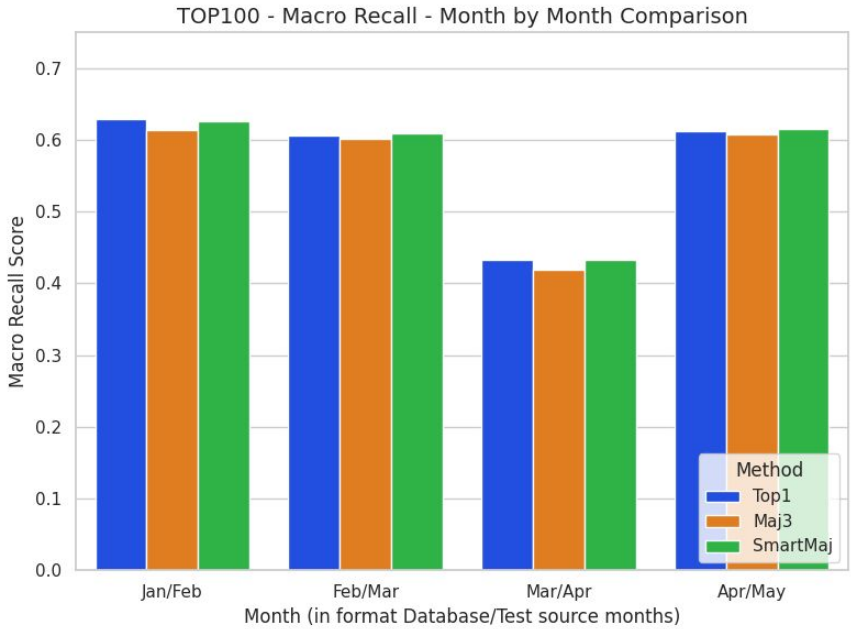

Transfer to ECH traffic

- We can build a database of non-ECH flows
  - Unless everyone uses ECH
- Assign a domain based on embeddings and their similarity
- Our next step and future work

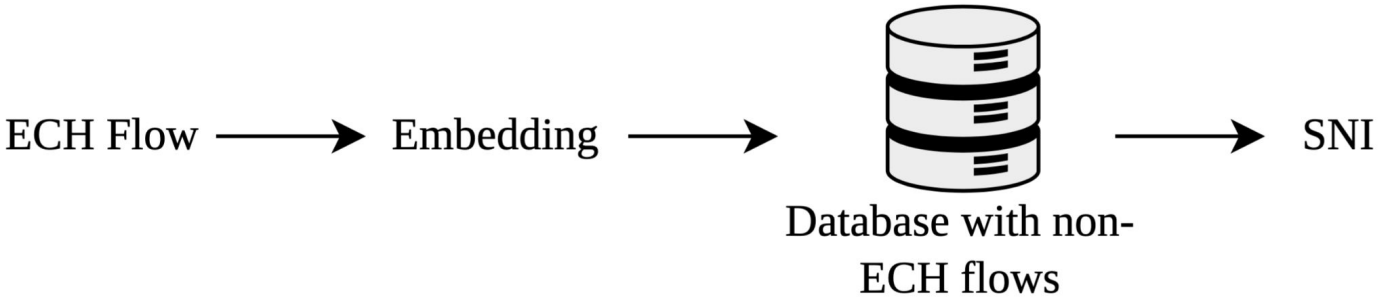

Findings & Results Interpretation

1. 30pktTCNET **can be used** to process new ECH (GREASE) data
2. Performance lowers with more domains
3. No clear winner emerged so far
4. CESNET-QUICEXT-25 **contains** some kind of **drift** in April 2025

Links for the mentioned papers and dataset

Thank you for your attention

Richard Plny, richard.plny@cesnet.cz

Jan Luxemburk, jan.luxemburk@cesnet.cz

Jaroslav Pesek, jaroslav.pesek@cesnet.cz

References

[1] Jonas Mücke, Konstantin Gasser, Thomas C. Schmidt, and Matthias Wählisch. 2025. Towards a Complete View of Encrypted Client Hello Deployments. In Proceedings of the ACM SIGCOMM 2025 Posters and Demos (ACM SIGCOMM Posters and Demos '25). Association for Computing Machinery, New York, NY, USA, 22–24. https://doi.org/10.1145/3744969.3748401

[2] J. Luxemburk, K. Hynek, R. Plný and T. Čejka, "Universal Embedding Function for Traffic Classification via QUIC Domain Recognition Pretraining: A Transfer Learning Success," in IEEE Transactions on Network and Service Management, vol. 23, pp. 1647-1663, 2026, doi: 10.1109/TNSM.2025.3642984.