Chair of Connected Mobility
TUM School of Computation, Information and Technology
Technische Universität München

Future Internet Protocols, Consolidation and Decentralization

David Guzman, Srivatsav Chenna

Chair of Connected Mobility

TUM School of Computation, Information and Technology (CIT)

Outline

1. Structure and Goals of the Seminar
   a. Seminar Structure and Grading Scheme (Reviews + Presentation)
   b. Learning Goals and Outcomes
2. Reading Papers
3. Reviews
   a. HotCRP Example
   b. Reviewing Process
   c. Writing Hints
4. Presentation Hints
5. Organizational Matters
6. Overview of Pre-selected Papers (Moodle)

1. Seminar Structure and Grading Scheme

- Regular conference paper publication process:
  - Call for Papers: Submission until specified date
  - Round of reviews by technical committee: decision to accept or reject
  - Presentation of accepted papers during conference sessions (“camera-ready”)
    - Followed by discussion and Q&A
- In this Seminar: Reviews + Presentation — ***Grading***
  - Pre-selected (published) papers on Secure Protocol Design
  - Each student reviews papers for the sessions — ***40%***
    - Focus on relevance for future work and discussion questions
  - Each student presents 1 paper **(~12-13 min + 8-7 min discussion)** — ***40%***
    - Include your own thoughts on the paper in 1 slide
  - Discussion participation throughout the seminar — ***20%***

1. Learning Goals and Outcomes

- Learn about latest research in the area of **Consolidation, Decentralization and Distributed Protocol Design**
  - Current problems, state-of-the-art methodology, open questions
- Acquire general skill set by taking roles of a reviewer and presenter
  - Scientific conference/workshop, peer reviewing, paper publication process
    - How are good research papers presented in written form?
    - Importance of reproducibility and artifacts
  - Critical thinking while reading/reviewing papers and listening to talks
    - Questioning content, identifying strengths/weaknesses and impact
    - Extract key results and main findings → takeaways
    - Presentation and discussion of results in concise manner for a talk

2. Reading Papers

Keshav's three-pass approach (https://doi.org/10.1145/1273445.1273458)

1. Bird's-eye view (5 – 10 minutes)
   - Read abstract, introduction
   - Check section headings and references
   - Read conclusion and for known references
2. Key points (1 hour)
   - Read the full paper, ignoring details (e.g., proofs, mathematical formulas, …)
   - Look at figures and mark important references
3. Full understanding (4 – 5 hours)
   - "Re-implement" the paper from the same assumptions **(or at least try to!)**
   - Challenge assumptions and details; "How would I have done this?"
   - Find implicit assumptions, missing related work and other issues (e.g., shoddy evaluation)

2. From Paper Understanding to Review Writing

https://www.ub.tum.de/en/bookmarklet (Access papers using the TUM license)

https://fabio.pierazzi.com/blog/2021/literature-review-systems-security/ (Finding papers)

https://www2.cs.uregina.ca/~pwlfong/CS499/reading-paper.pdf (Reading CS papers)

https://cseweb.ucsd.edu//~wgg/CSE210/howtoread.html (Reading engineering papers)

https://gernot-heiser.org/benchmarking-crimes.html (Systems Benchmarking Crimes)

https://cs.brown.edu/~sk/Memos/Paper-Reviews/ (Writing reviews for technical papers)

https://people.inf.ethz.ch/troscoe/pubs/review-writing.pdf (Writing a systems review)

https://dl.acm.org/doi/10.1145/1519103.1519122 (How to not write a review)

3. HotCRP Example https://hotcrp.com

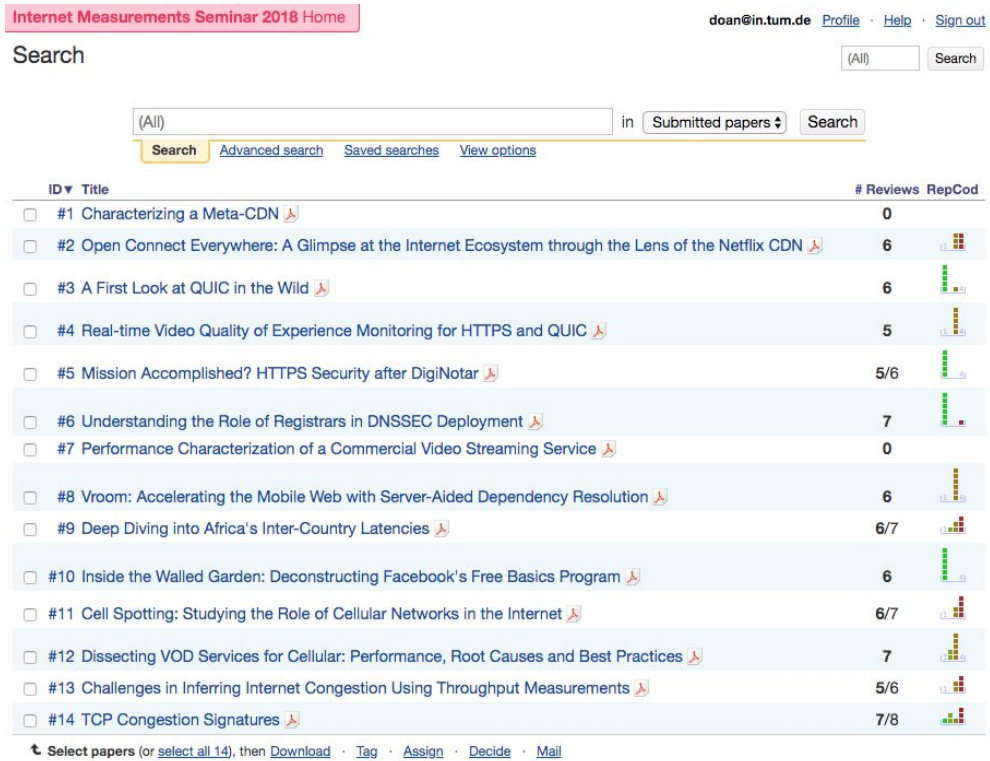

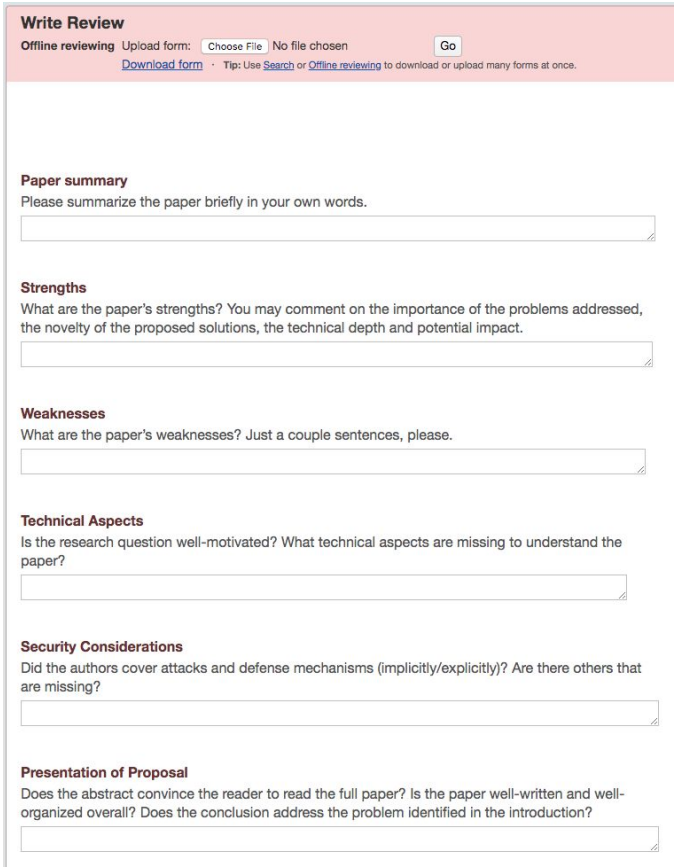

3. Reviewing Process

- We invite each of you to our HotCRP instance.
- Check email address (login) in invitation email.
- Registration form: affiliation/collaborators can be anything (e.g., TUM), unused field
- Presented papers will be imported to HotCRP for reviewing.
- Review form provides prompts to assist in commenting on the papers.
- Your 2 reviews are due until the presentation date **if not presenting** on that day.
  - i.e., before a session you either review (X)OR prepare a presentation.
- Overall each of you will write 4 reviews throughout the seminar.

3. Writing and Content Hints

- For Reviews (and Presentations!):
  - Use tools like Google Scholar, Connected Papers, ResearchRabbit, NotebookLM or Undermind to do a literature survey.
  - Find the **context of the paper, the state of the art and the research gap.**
    - E.g., what else have people tried, do newer approaches exist?
    - Beware that AI-based tools might hallucinate.
  - Based on that context, argue about the **impact** that the paper has **on the state of the art and on future work.**
  - **Since above tools do a lot of the analysis of strengths and weaknesses automatically for you, our review grading will focus on you conveying the key ideas of the paper and the impact the paper had.**
- For in-depth questions and discussion items:
  - Avoid generic questions such as *“Can we apply AI/ML?”*
  - Try to ask questions that, e.g., propose changing the design or configuration of the study and briefly elaborate on why

NotebookLM

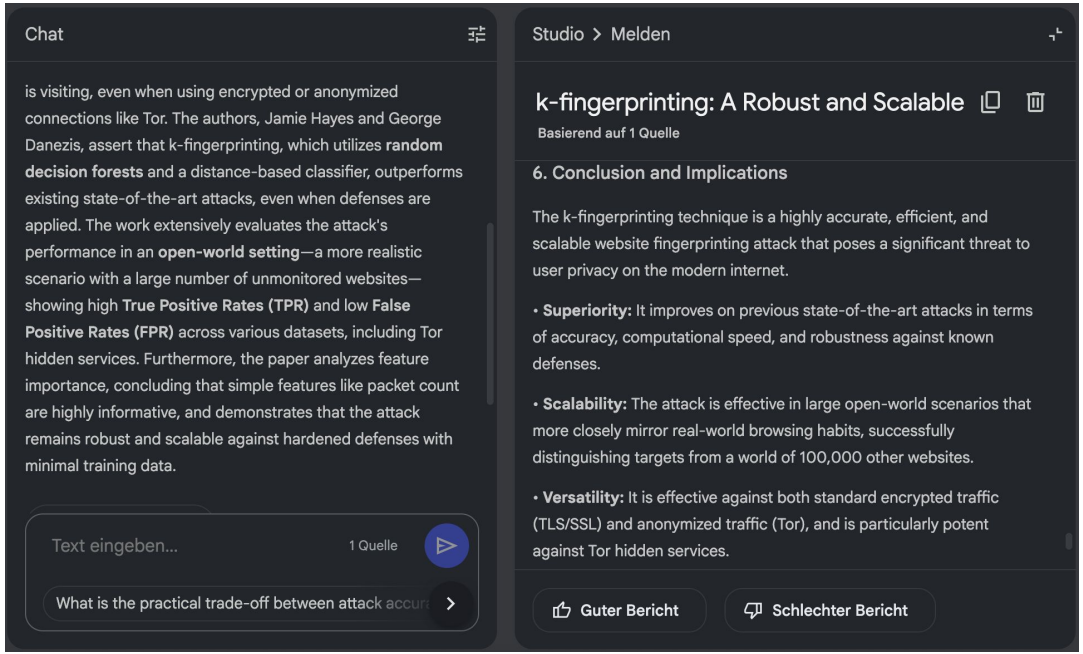
Chat
is visiting, even when using encrypted or anonymized connections like Tor. The authors, Jamie Hayes and George Danezis, assert that k-fingerprinting, which utilizes **random decision forests** and a distance-based classifier, outperforms existing state-of-the-art attacks, even when defenses are applied. The work extensively evaluates the attack's performance in an **open-world setting**—a more realistic scenario with a large number of unmonitored websites—showing high **True Positive Rates (TPR)** and low **False Positive Rates (FPR)** across various datasets, including Tor hidden services. Furthermore, the paper analyzes feature importance, concluding that simple features like packet count are highly informative, and demonstrates that the attack remains robust and scalable against hardened defenses with minimal training data.
Text eingeben...
1 Quelle
What is the practical trade-off between attack accur
Studio > Melden
k-fingerprinting: A Robust and Scalable
Basierend auf 1 Quelle
6. Conclusion and Implications
The k-fingerprinting technique is a highly accurate, efficient, and scalable website fingerprinting attack that poses a significant threat to user privacy on the modern internet.
• **Superiority:** It improves on previous state-of-the-art attacks in terms of accuracy, computational speed, and robustness against known defenses.
• **Scalability:** The attack is effective in large open-world scenarios that more closely mirror real-world browsing habits, successfully distinguishing targets from a world of 100,000 other websites.
• **Versatility:** It is effective against both standard encrypted traffic (TLS/SSL) and anonymized traffic (Tor), and is particularly potent against Tor hidden services.
Guter Bericht
Schlechter Bericht

Connected Papers

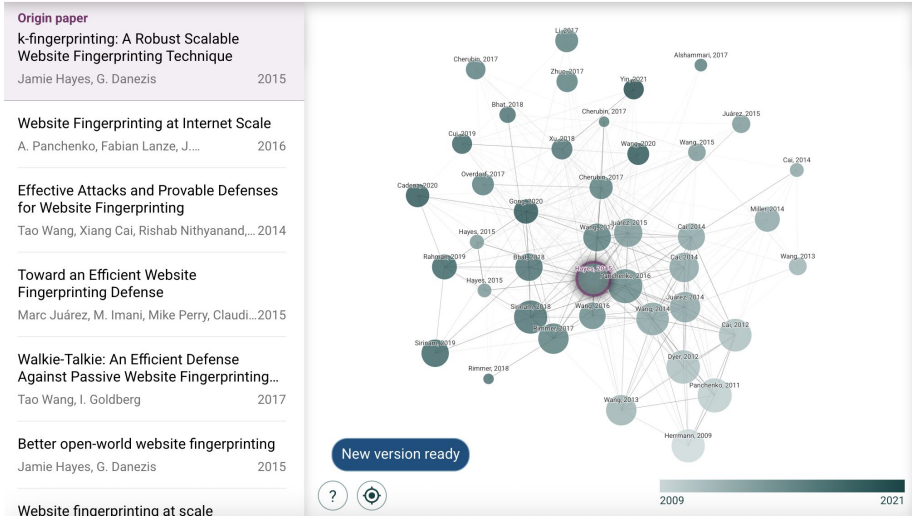

Undermind

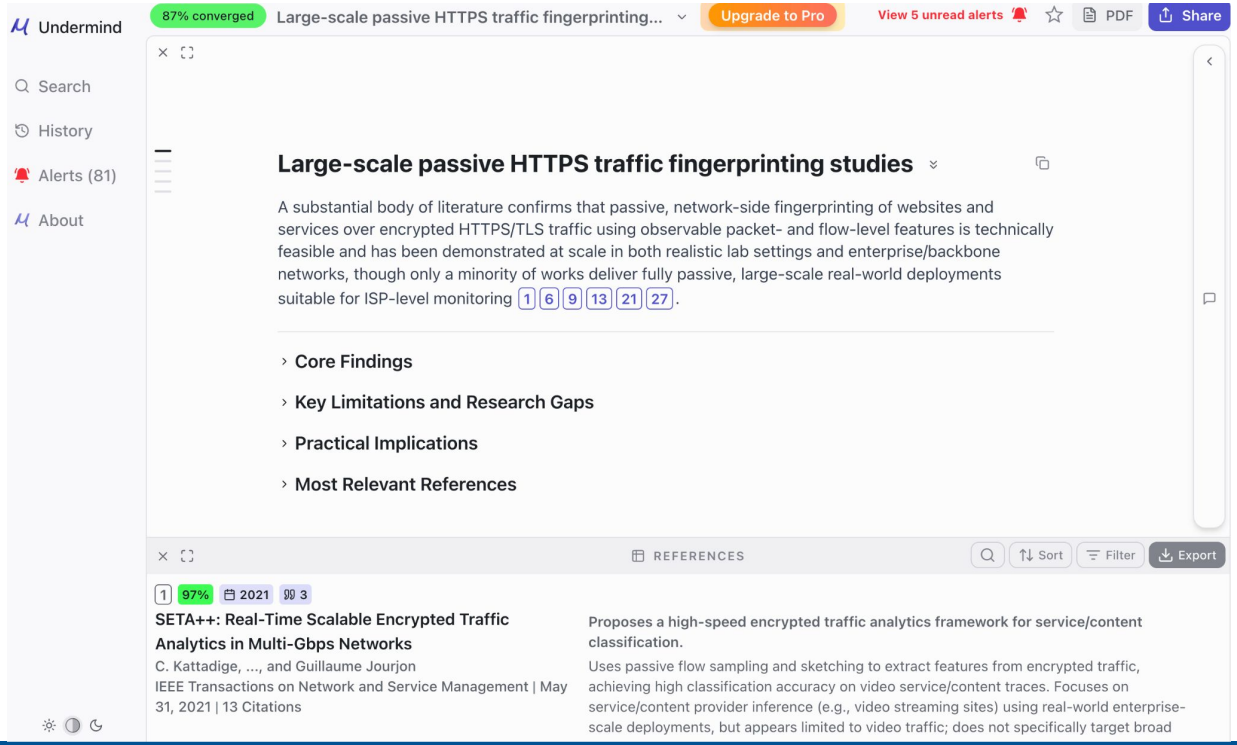
Undermind
87% converged
Large-scale passive HTTPS traffic fingerprinting...
Upgrade to Pro
View 5 unread alerts
PDF
Share
Search
History
Alerts (81)
About
Large-scale passive HTTPS traffic fingerprinting studies
A substantial body of literature confirms that passive, network-side fingerprinting of websites and services over encrypted HTTPS/TLS traffic using observable packet- and flow-level features is technically feasible and has been demonstrated at scale in both realistic lab settings and enterprise/backbone networks, though only a minority of works deliver fully passive, large-scale real-world deployments suitable for ISP-level monitoring 1 6 9 13 21 27.
Core Findings
Key Limitations and Research Gaps
Practical Implications
Most Relevant References
REFERENCES
Sort
Filter
Export
1
97%
2021
3
SETA++: Real-Time Scalable Encrypted Traffic Analytics in Multi-Gbps Networks
C. Kattadige, ..., and Guillaume Jourjon
IEEE Transactions on Network and Service Management | May 31, 2021 | 13 Citations
Proposes a high-speed encrypted traffic analytics framework for service/content classification.
Uses passive flow sampling and sketching to extract features from encrypted traffic, achieving high classification accuracy on video service/content traces. Focuses on service/content provider inference (e.g., video streaming sites) using real-world enterprise-scale deployments, but appears limited to video traffic; does not specifically target broad

3. More Writing and Content Hints

- Be **concise** and **specific**; make sure your arguments are **logical** and **understandable**.
- Use **paragraphs** to make reading easier. Alternatively use bullet points (but please write proper sentences).
- **Avoid complex** sentences, sentences **without meaning** and **hyperboles** (e.g., “greatest”).
- Try to be **constructive**: when criticizing something, try to understand the author’s intent and give concrete steps for improvement.
- We will check whether you used ChatGPT for writing and will run reviews through a plagiarism checker!
- We will not provide example reviews. You can take a look at past reviews from WWW:
  - https://openreview.net/group?id=ACM.org/TheWebConf/2025/Conference
  - https://openreview.net/group?id=ACM.org/TheWebConf/2024/Conference
  - NOTE: Most of these reviews are not going to be good examples because reviewers tend to have less time than you do in this course.

4. Presentation Hints

- Presentations typically follow the structure of the paper
  - Motivate the subject **(~15% of time)**
  - Explain methodology or system design **(~40% of time)**
    - Keep it simple: don’t lose yourself in details; add background if needed
  - Present and discuss **main/most interesting results** **(~30% of time)**
    - Usually summarized in abstract
  - Summarize main findings/conclusions and describe outlook **(~10% of time)**
  - Add your own thoughts on the paper in 1 slide **(~5% of time)**
    - e.g., strengths, weaknesses, limitations
- Practice your presentation
  - Get feedback from your fellow students
  - Be aware of the **time limit of ~12-13 min**

4. More Presentation Hints

No outline or table of contents is necessary unless you deviate from this structure quite a bit!

(see https://www.microsoft.com/en-us/research/academic-program/give-great-research-talk/ – “Research is Communication”)

1. Motivation (see https://cs.stanford.edu/people/widom/paper-writing.html) **(~3 minutes)**
   - “Always ask why” (but keep it short)
   - What is the problem? (problem statement, *not* research questions) → Why is it interesting and important? (context) → Why is it hard? (at a high level) → Why hasn't it been solved before? (existing state of the art) → What are the key components and results? (research questions, subproblems)
2. Methodology/System Design **(~8 minutes)**
   - Add background here if needed: explain things that are not in the paper but you yourself needed to understand the paper
   - Datasets, data preparation, measurement setup, system design from a top down perspective
   - Depending on the paper: reduce the granularity for the presentation (omit technical details and math unless it is key to understanding)
   - You can also mention trade-offs that you believe the design to have made here
3. Main Results/System Evaluation **(~6 minutes)**
   - You can take a different “logical flow” than the one used by the paper
   - Key insights and takeaways that the authors presented (and how they got there based on the methodology or how the evaluation helps answer a subproblem)
   - Crop, copy and edit figures from the paper if needed
4. Summary and Conclusion **(~2 minutes)**
   - What problem did the paper solve and what are the most important contributions?
   - What are the next steps the authors propose?
5. Own Thoughts **(~1 minute)**
   - Strengths, weaknesses and limitations of the [approach | measurements | system | techniques] from your own point of view

5. Organizational Matters

- Paper selection
  - List of seminar papers in Moodle
  - Send us your preferences by **May 6th 23:59** via e-mail:
    - **Top 7 papers in descending order**
  - Your top-most choice indicates the paper you would like to present most.
  - We will try to assign reviews based on your preferences. But the overall goal is that the papers that are presented also get reviewed.
  - We will assign the papers and presentation dates by October 27, and announce them via Moodle
- We want to figure out if and how you are using AI tools and whether they actually helped you.
  - Thus, each presenter SHOULD create one slide about interesting findings when using AI tools. If you did not use any, you do not have to do this.
  - We will make time for briefly discussing this at the end of each session.

5. Organizational Matters (contd.)

- Seminar sessions will take place roughly every two weeks (to provide enough time for reading and preparation of papers) in room **01.13.010** or as an exception on BBB https://bbb.cit.tum.de/rooms/bii-0fp-4ll-uqb/join
- Sessions start at 12 p.m.
- Recommended (but not necessary!) background reading: Future Internet Architectures
- Next dates:
  - Preferences – Send by May 6 **(check Moodle!)**
  - Introduction – Wednesday April 22, 12:00 - 14:00
  - Session 1 – Wednesday May 20, 12:00 - 14:00
  - Session 2 – Wednesday June 3, 12:00 - 14:00
  - Session 3 – Wednesday June 17, 12:00 - 14:00
  - Grading – Wednesday August 1

6. Papers

Algorithms

Mercury: Fast Transaction Broadcast in High Performance Blockchain Systems. Infocom’23

Strategic Latency Reduction in Blockchain Peer-to-Peer Networks. PAM’23

Kadcast: A Structured Approach to Broadcast in Blockchain Networks. (Workshop)’2019

Perigee: Efficient Peer-to-Peer Network Design for Blockchains. PODC’2020

CommiTEE: An Efficient and Secure Commit-Chain Protocol using TEEs

Falcon: Advancing Asynchronous BFT Consensus for Lower Latency and Enhanced Throughput (VLDB 25)

6. Papers

Decentralizing Systems

Virtual Stars, Real Fans: Understanding the VTuber Ecosystem. WWW'24

Design and evaluation of ipfs: A storage layer for the decentralized web. Sigcomm'22

Looking AT the Blue Skies of Bluesky IMC'24

Bootstrapping Social Networks: Lessons from Bluesky. ICWSM'25

ReP2P Matrix: Decentralized Relays to Improve Reliability and Performance of Peer-to-Peer Matrix. CoNEXT'24

Towards a Decentralized Internet Namespace. CoNEXT'24

A Prototype for Privacy-preserving and Compliant Offline CBDC Transactions (BRAINS 25)

A Secure and Flexible Blockchain-Based Offline Payment Protocol (IEEE Transactions on Computers)

An Empirical Analysis of the Nostr Social Network: Decentralization, Availability, and Replication Overhead (CoNEXT'25)

6. Papers

What is happening in Decentralized Systems? (Pros and Cons)

Fediverse Migrations: A Study of User Account Portability on the Mastodon Social Network. IMC'24

The Cloud Strikes Back : Investigating the Decentralization of IPFS. 391–405. https://doi.org/10.1145/3618257.3624797

Piecing Together the Jigsaw Puzzle of Transactions on Heterogeneous Blockchain Networks. Sigmetrics'24

A flash(bot) in the pan: measuring maximal extractable value in private pools IMC'2022

Communication Cost for Permissionless Distributed Consensus at Internet Scale. CoNEXT'24

Ethereum Price Prediction Employing Large Language Models for Short-term and Few-shot Forecasting (BRAINS 25)

Forensic Analysis of the Libra Memecoin Incident (BRAINS 25)

Sealed-Bid Auctions via TEE-Backed Confidential Compute

A Systematisation of Knowledge: Connecting European Digital Identities with Web3 (Blockchain 24)

Your turn!

Tell us something about you :)

E.g., Name, study course, interests