Privacy Proofs for Anonymous Communication Networks

Christoph Coijanovic Daniel Schadt Thorsten Strufe

KASTEL Security Research Labs, Karlsruhe Institute of Technology

Abstract

Today's online communications expose highly personal *metadata*, such as who is communicating with whom, to service providers and other potential observers. Anonymous Communication Networks (ACNs) use cryptographic building blocks to hide metadata in communications. In order to gain confidence in the security guarantees that an ACN can provide, various formalization approaches for metadata protection and proof frameworks have been proposed [KBS20; Kuh+19; PH10; Bac+16]. In this talk, we give an overview of existing proof techniques for ACNs and discuss our findings from applying these techniques to concrete systems.

Motivation *Metadata* in online communication, such as who someone is communicating with, is often used to marginalise, discriminate against, or even prosecute people [MMM16]. In response, anonymous communication has been an active research topic for decades, with many systems proposed [SG24]. These anonymous communication networks (ACNs) combine cryptographic building blocks with additional techniques such as cover traffic to hide communication metadata.

As ACNs are complex systems, potentially deployed in high-stakes environments, it is important that their security and privacy guarantees are rigorously proven. If a formal proof of privacy exists for a given system, one can a) better understand the exact metadata that the system aims to protect, and b) be confident that this protection is actually achieved. Within the anonymous communication community, there are different approaches to formalising and proving privacy.

Existing Approaches To prove privacy in the systems we have designed in the past, we have followed two basic strategies, which we will contrast in this talk.

For some systems ([Gab+21; CWS23; Coi+24] and upcoming work), we build on Kuhn et al.'s framework of *privacy notions*. Through this framework, Kuhn et al. propose a game-based approach to modeling metadata disclosure in ACNs. Similar to the IND-CPA game in cryptography, the game is played

between a challenger and an adversary. The adversary is allowed to choose two alternative communication patterns for clients and then has to decide, based on the disclosed metadata, which of these patterns the challenger has chosen to execute in the ACN. Note that this game is independent of both the concrete ACN and the metadata to be hidden. Concrete ACNs are mapped to "protocol models", which simplify the protocol to a description of the metadata that would be revealed to an adversary observing an implementation. The privacy goals of the ACN are applied to the game by restricting how the alternative communication patterns may differ from each other; any metadata that the ACN does not wish to hide must be identical in both patterns, so that it does not help the adversary to distinguish between them. Because of the universal nature of this approach, it can be used to analyze a wide variety of ACN designs and compare their privacy guarantees.

For other systems ([Sch+24] and upcoming work) we have relied on lower-level properties tailored to the specific protocol functionality. These properties are also formalized by an indistinguishability game, but restrict the adversary to one possible attack vector. For example, one property ("Layer Unlinkability") requires that an adversary observing an anonymization server cannot link incoming and outgoing packets. Compared to Kuhn et al.'s privacy notions, where the adversary inputs communication patterns and receives all disclosed information through the protocol model as output, the properties give the adversary direct access to protocol functions: In the case of Layer Unlinkability, the adversary can select two alternative inputs to the packet processing function on the honest server and receives back the function's output from one of the inputs. Taken together, all targeted properties define the privacy provided by the protocol. With this approach, proofs are generally simpler because the scope of each proof is narrower. However, comparisons between different families of protocols are difficult because the properties only apply to protocols in the same family.

Research and Talk In our research group, we have experience both in formalizing notions of privacy and in using these formalizations to prove the privacy of ACNs we have designed. In this talk, we will discuss our findings from proving the privacy of four published systems [Coi+24; CWS23; Sch+24; Gab+21] and two upcoming ones. We intend to focus on lessons learned and provide guidance on the different approaches to proving privacy guarantees for anonymous communication networks.

References

[PH10] Andreas Pfitzmann and Marit Hansen. "A terminology for talking about privacy by data minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management". In: 2010.

[Bac+16] Michael Backes et al. "AnoA: A Framework for Analyzing Anonymous Communication Protocols". In: *J. Priv. Confidentiality* 7.2 (2016).

[MMM16] Jonathan R. Mayer, Patrick Mutchler, and John C. Mitchell. "Evaluating the privacy properties of telephone metadata". In: *Proc. Natl. Acad. Sci. USA* 113.20 (2016), pp. 5536–5541.

[Kuh+19] Christiane Kuhn et al. "On Privacy Notions in Anonymous Communication". In: *Proc. Priv. Enhancing Technol.* (2019).

[KBS20] Christiane Kuhn, Martin Beck, and Thorsten Strufe. "Breaking and (Partially) Fixing Provably Secure Onion Routing". In: *IEEE Symposium on Security and Privacy (SP)*. 2020.

[Gab+21] Sarah Abdelwahab Gaballah et al. "2PPS - Publish/Subscribe with Provable Privacy". In: *40th International Symposium on Reliable Distributed Systems, SRDS 2021*. IEEE, 2021.

[CWS23] Christoph Coijanovic, Christiane Weis, and Thorsten Strufe. "Panini - Anonymous Anycast and an Instantiation". In: *Computer Security - ESORICS 2023 - 28th European Symposium on Research in Computer Security*. Vol. 14345. Lecture Notes in Computer Science. Springer, 2023, pp. 193–211.

[Coi+24] Christoph Coijanovic et al. "Pirates: Anonymous Group Calls over Fully Untrusted Infrastructure". In: *Information Security and Privacy - 29th Australasian Conference, ACISP 2024*. Vol. 14897. Lecture Notes in Computer Science. Springer, 2024, pp. 193–212.

[SG24] Sajin Sasy and Ian Goldberg. "SoK: Metadata-Protecting Communication Systems". In: *Proc. Priv. Enhancing Technol.* 2024.1 (2024), pp. 509–524.

[Sch+24] Daniel Schadt et al. "PolySphinx: Extending the Sphinx Mix Format With Better Multicast Support". In: *IEEE Symposium on Security and Privacy (SP)*. 2024.