AnoA: A Framework For Analyzing Anonymous Communication Protocols
Unified Definitions and Analyses of Anonymity Properties

Michael Backes
Saarland University and MPI-SWS

Aniket Kate
MMCI, Saarland University

Praveen Manoharan
Saarland University

Sebastian Meiser
Saarland University

Esfandiar Mohammadi
Saarland University

I. Introduction

Protecting individuals' privacy in online communications has become a challenge of paramount importance. A wide variety of privacy enhancing technologies, comprising many different approaches, have been proposed to solve this problem. Privacy enhancing technologies, such as anonymous communication (AC) protocols, typically protect users' privacy by anonymizing their communication over the Internet. Employing AC protocols has become increasingly popular over the last decade. This popularity is exemplified by the success of the Tor network [1].

There has been significant previous work on analyzing the anonymity provided by various AC protocols such as dining cryptographers network (DC-net) [2], mix network (Mixnet) [3], and onion routing (e.g., Tor) [4]. (See [5] and references therein.) However, most of the previous work only considers a single anonymity property for a particular AC protocol under a specific adversary scenario. So far, there is no framework that is both expressive enough to unify and compare relevant anonymity notions (such as sender anonymity, sender unlinkability, and relationship anonymity), and that is also well suited for analyzing complex cryptographic protocols.

Previous frameworks such as [6] only guarantee anonymity for a symbolic abstraction of the AC, not for its cryptographic realization. Moreover, while some existing work, such as [5], considers an adversary with access to *a priori* probabilities for the behavior of users, there is still no framework that is capable of dealing with an adversary that has arbitrary auxiliary information about user behavior.

II. Contributions

In this work, we make three contributions to the field of anonymity analysis.

As a first contribution, we present the novel anonymity analysis framework AnoA. In AnoA we define and analyze anonymity properties of AC protocols. AnoA is compatible with simulation-based composability frameworks, such as Universal Composability (UC), the IITM model, or Reactive Simulatability (RSIM). In particular, for all protocols that are securely abstracted by an ideal functionality [7], [8], our definitions allow an analysis of these protocols in a symbolic manner.

As a second contribution, we formalize the well-established notions of sender anonymity, (sender) unlinkability, and relationship anonymity in our framework. We discuss why our anonymity definitions accurately capture these notions. Moreover, we show the relations between our formalizations of sender anonymity, (sender) unlinkability, and relationship anonymity.

As a third contribution, we apply our framework to the most successful AC protocol: Tor. We give quantitative results for obtaining anonymity according to all of above anonymity notions by using Tor [1].

III. Overview

We give a brief description of our work. The formal definitions, theorems and proofs can be found in our technical report at http://www.infsec.cs.uni-saarland.de/~meiser/paper/anoa.html.

A. *The* AnoA *Framework*

Our anonymity framework AnoA is based on a novel generalization of computational differential privacy (IND-CDP), a notion introduced by Mironov et al. [9]. Informally, differential privacy of a mechanism guarantees that this mechanism does not leak any information about a single user – even to an adversary that has auxiliary information about the rest of the user base. IND-CDP compares two *adjacent* input tables, i.e., input tables that differ in one row. The definition basically states that no ppt adversary should be able to determine which of these input tables was used.

For anonymity properties of AC protocols, such a notion of adjacency is too strong. One of the main objectives of an AC protocol is communication: delivering the sender's message to the recipient. However, if these messages carry information about the sender, a curious recipient can determine the sender.

```
Upon message(input, $D_0, D_1$) (only once)
  compute $(D'_0, D'_1) \leftarrow \alpha(D_0, D_1)$
  if $(D'_0, D'_1) \neq \bot$ then
    run $\mathcal{P}$ on the input table $D'_b$ and forwards all
    messages that are sent by $\mathcal{P}$ to $\mathcal{A}$ and vice versa.
```

Figure 1. The challenger $\mathrm{CH}_b(\mathcal{P}, \alpha)$ for the adjacency function α

Additionally, to specify different variants of anonymity (e.g. sender unlinkability and relationship anonymity), we want to give purpose-specific descriptions of the inputs we consider adjacent.

These observations lead to our new notion of α-IND-CDP that allows a unified specification of anonymity properties based on adjacency functions α:

For analyzing a protocol $\mathcal{P}$, we define a challenger CH_b that expects two input tables from a ppt adversary $\mathcal{A}$. The challenger CH_b chooses one of them and successively sends one row after the other to the protocol $\mathcal{P}$. It is straightforward to construct a wrapper for $\mathcal{P}$ that translates input tables to the expected input of the protocol.

Definition 1 ((ε, δ)-α-IND-CDP). *Let* CH_b *be the challenger from Figure 1. The protocol* $\mathcal{P}$ *is* (ε, δ)-α-IND-CDP *for* α, *where* $\varepsilon \geq 0$ *and* $0 \leq \delta \leq 1$, *if for all ppt-adversaries* $\mathcal{A}$:

$$\begin{aligned} &\Pr[b = 0 : b \leftarrow \mathcal{A}^{\mathrm{CH}_0(\mathcal{P}, \alpha)}] \\ &\leq e^{\varepsilon} \cdot \Pr[b = 0 : b \leftarrow \mathcal{A}^{\mathrm{CH}_1(\mathcal{P}, \alpha)}] + \delta \end{aligned}$$

We stress that our protocol model is generic enough to capture multi-party protocols in classical simulation-based composability frameworks, such as the UC, the IITM or the RSIM framework. In particular, our framework is strong enough to guarantee anonymity properties for an AC protocol as soon as these properties have been proven for a securely realized ideal functionality. In the full version we formally prove this for UC.

B. Studying our Anonymity Definitions

We give definitions for the adjacency functions that capture the notions of sender anonymity (α_{SA}), sender unlinkability (α_{UL}) and relationship anonymity (α_{Rel}). We show that our definitions accurately capture these anonymity notions, and show that our definitions are equivalent to the definitions from the literature.[1]

The unified approach of defining anonymity notions in AnoA allows for a formal comparison of these notions. We use this fact to give formal proofs for the relations between each of the notions. Figure 2 illustrates our results.

$$\begin{tikzcd}
\alpha_{\mathrm{UL}} \arrow[r, Rightarrow, bend right, "/"{marking}] & \alpha_{\mathrm{SA}} \arrow[l, Rightarrow, bend right] \arrow[r, Rightarrow, bend left] & \alpha_{\mathrm{Rel}} \arrow[l, Rightarrow, bend left, "/"{marking}]
\end{tikzcd}$$

Figure 2. The relations between our anonymity definitions

C. Tor Analysis

We leverage previous results that securely abstract Tor as an ideal functionality (in the UC framework) [8]. Then, we illustrate that proving sender anonymity, sender unlinkability and relationship anonymity against passive adversaries boils down to a combinatoric analysis based on the number of corrupted nodes in the network.

Since the underlying cryptographic model does not capture system-level attacks, we model known system-level attacks, such as website fingerprinting and traffic correlation, as over-approximation in the ideal functionality.

In addition, we discuss a known counter-measure for Tor's high sensitivity to compromised nodes: the entry guards mechanism. We show that using entry guards dramatically reduces the adversary's success probability and why this is the case.

References

[1] "The Tor Project," https://www.torproject.org/, 2003, accessed Feb 2013.

[2] D. Chaum, "The dining cryptographers problem: Unconditional sender and recipient untraceability," *J. Cryptology*, vol. 1, no. 1, pp. 65–75, 1988.

[3] ——, "Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms," *Communications of the ACM*, vol. 4, no. 2, pp. 84–88, 1981.

[4] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The Second-Generation Onion Router," in *Proc. 13th USENIX Security Symposium (USENIX)*, 2004, pp. 303–320.

[5] J. Feigenbaum, A. Johnson, and P. F. Syverson, "Probabilistic analysis of onion routing in a black-box model," *ACM Trans. Inf. Syst. Secur.*, vol. 15, no. 3, p. 14, 2012.

[6] D. Hughes and V. Shmatikov, "Information hiding, anonymity and privacy: a modular approach," *Journal of Computer Security*, vol. 12, no. 1, pp. 3–36, 2004.

[7] G. Danezis and I. Goldberg, "Sphinx: A Compact and Provably Secure Mix Format," in *Proc. 30th IEEE Symposium on Security & Privacy*, 2009, pp. 269–282.

[8] M. Backes, I. Goldberg, A. Kate, and E. Mohammadi, "Provably secure and practical onion routing," in *Proc. 26st IEEE Symposium on Computer Security Foundations (CSF)*, 2012, pp. 369–385.

[9] I. Mironov, O. Pandey, O. Reingold, and S. P. Vadhan, "Computational differential privacy," in *CRYPTO*, 2009, pp. 126–142.

[10] A. Pfitzmann and M. Hansen, "A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management," http://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf, Aug. 2010, v0.34.

[1] We particularly consider the definitions from the seminal work by Pfitzmann and Hansen [10].