

Anonymity Trilemma: Strong Anonymity, Low Bandwidth Overhead, Low Latency—Choose Two

Debajyoti Das
Purdue University, USA
das48@purdue.edu

Sebastian Meiser
University College London, UK
s.meiser@ucl.ac.uk

Esfandiar Mohammadi
ETH Zurich, Switzerland
mohammadi@inf.ethz.ch

Aniket Kate
Purdue University, USA
aniket@purdue.edu

Abstract—This work investigates the fundamental constraints of anonymous communication (AC) protocols. We analyze the relationship between bandwidth overhead, latency overhead, and sender anonymity or recipient anonymity against the global passive (network-level) adversary. We confirm the trilemma that an AC protocol can only achieve two out of the following three properties: strong anonymity (i.e., anonymity up to a negligible chance), low bandwidth overhead, and low latency overhead.

We further study anonymity against a stronger global passive adversary that can additionally passively compromise some of the AC protocol nodes. For a given number of compromised nodes, we derive necessary constraints between bandwidth and latency overhead whose violation make it impossible for an AC protocol to achieve strong anonymity. We analyze prominent AC protocols from the literature and depict to which extent those satisfy our necessary constraints. Our fundamental necessary constraints offer a guideline not only for improving existing AC systems but also for designing novel AC protocols with non-traditional bandwidth and latency overhead choices.

I. INTRODUCTION

Millions of users from all over the world employ anonymous communication networks, such as Tor [1], to protect their privacy over the Internet. The design choice made by the Tor network to keep the latency and bandwidth overheads small has made it highly attractive to its geographically diverse user-base. However, over the last decade, the academic literature [2]–[8] has demonstrated Tor’s vulnerability to a variety of traffic correlation attacks. In fact, Tor also has been successfully attacked in practice [9].

It is widely accepted that low-latency low-bandwidth overhead of anonymous communication (AC) protocols, such as Tor [10], can only provide a weak form of anonymity [11]. In the anonymity literature, several AC protocols were able to overcome this security barrier to provide a stronger anonymity guarantee (cryptographic indistinguishability based anonymity [12], [13]) by either increasing the latency overhead or the bandwidth overhead. In particular, high-latency approaches (such as threshold mix networks [14]) can ensure strong anonymity by introducing significant communication delays for users messages, while high-bandwidth approaches (such as Dining Cryptographers network [15] and its extensions [16]–[18]) can provide strong anonymity by adding copious noise (or dummy) messages.

There have been a few efforts to propose hybrid approaches [19]–[24] that try to provide anonymity by simultaneously introducing latency and bandwidth overhead. However,

it is not clear how to balance such system parameters to ensure strong anonymity while preserving practical performance.

In general, in the last 35 years a significant amount of research efforts have been put towards constructing novel AC protocols, deploying them, and attacking real-world AC networks. However, unlike other security fields such as cryptography, our understanding regarding the fundamental limits and requirements of AC protocols remains limited. This work takes some important steps towards answering fundamental question associated with anonymous communication. “Can we prove that strong anonymity cannot be achieved without introducing large latency or bandwidth overhead? When we wish to introduce the latency and bandwidth overheads simultaneously, do we know the overhead range values that still fall short at providing stronger anonymity?”

Our Contribution. We confirm a previously conjectured [24], [25] relationship between bandwidth overhead, latency overhead and anonymity. We find that there are fundamental bounds on sender and recipient anonymity properties [12], [13], [26], [27] of a protocol that directly depend on the introduced bandwidth and latency overheads.

This work presents a generic model of AC protocols using petri nets [28], [29] such that different instantiations of this model will represent different AC protocols, covering most practical AC systems in the literature. We derive *upper* bounds on anonymity as functions of bandwidth overhead and latency overhead, against two prominent adversary classes: global passive network-level adversaries and strictly stronger adversaries that additionally (passively) compromise some protocol parties (e.g., relays in case of Tor). These bounds constitute necessary constraints for anonymity. Naturally, the constraints are valid against any stronger adversary class as well.

For both adversary classes, we analyze two different user distributions (i.e., distributions that determine at which time or rate users of the AC protocol send messages): (i) synchronized user distributions, where users globally synchronize their messages, and (ii) unsynchronized user distributions, where each user locally decides when to send his messages independent of other users.

We analyze the trade-off between latency overhead and bandwidth overhead required to achieve *strong anonymity*, i.e., anonymity up to a negligible (in a security parameter η) chance of failure. For any AC protocol where only a fraction

of $\beta \in [0, 1]$ users send noise messages per communication round, and where messages can only remain in the network for $\ell \geq 0$ communication rounds, we find that against a global network-level adversary no protocol can achieve strong anonymity if $2\beta\ell < 1 - 1/\text{poly}(\eta)$ even when all the protocol parties are honest. In the case where a strictly stronger adversary additionally passively compromises c (out of K) protocol parties, we show that strong anonymity is impossible if $2(\ell - c)\beta < 1 - 1/\text{poly}(\eta)$ (for $c < \ell$), or $2\beta\ell < 1 - 1/\text{poly}(\eta)$ and $\ell \in \mathcal{O}(1)$ (for $c \geq \ell$).

We also assess the practical impact of our results by analyzing prominent AC protocols. Our impossibility results naturally only offer necessary constraints for anonymity, but *not* sufficient conditions for the AC protocol. However, these necessary constraints for sender and recipient anonymity are crucial for understanding bi-directional anonymous communication. In fact, we find that several AC protocols in the literature are asymptotically close to the suggested constraints. Moreover, designers of new AC protocols can use our necessary constraints as guidelines for avoiding bad trade-off between latency and bandwidth-overhead.

II. OVERVIEW

A. Formalization and Adversary Model

AC Protocols as Petri Nets. We define a view of AC protocols as petri nets [28]–[30], i.e., as graphs with two types of labeled nodes: *places*, that store colored tokens, and *transitions*, that define how these tokens are sent over the graph. In our case, each colored token represents a message, places are the protocol parties that can receive, hold and send messages, and transitions describe how parties exchange and relay messages. Our model captures all AC protocols under the assumption that messages are transmitted directly, i.e., in order for Bob to receive a message from Alice, Alice has to send the message and the message (albeit relayed, delayed and cryptographically modified) eventually has to reach Bob. While this requirement may sound strict, as elaborated in Section IV-B, we effectively only exclude few esoteric protocols.

User Distributions, Communication Rounds, Bandwidth Overhead, and Latency. We consider two types of *user distributions*. In the first user distribution (*synchronized*) N users send their messages in exactly N rounds (see Figure 1 for notations). Per round, exactly one user sends a message. The protocol decides which users send noise messages in each round. In the second user distribution (*unsynchronized*) each user independently decides whether to send a message in a round using a coin flip, with a success probability p .

The model considers synchronous communication *rounds* as in [16], [17], [31], [32]. We model latency overhead ℓ as the number of rounds a message can be delayed by the protocol before being delivered. We formalize bandwidth overhead β as the number of noise messages per user that the protocol can create in every round, i.e., the dummy message rate.

Our two types of user distributions cover a large array of possible scenarios. Results for our user distributions imply

ℓ	Latency overhead for every message
β	Bandwidth overhead for every user per round
p	Probability to send a message per user per round
K	Number of (internal) protocol parties
c	Number of compromised protocol parties
N	Number of online users (that may send messages)
δ	Adversarial advantage in the anonymity game
Π	A protocol. $\Pi \in \mathcal{M}$: Π is within our model
η	The security parameter
ϵ	A (very small, but non-negligible) function

Fig. 1. Notation

results for similar distributions, if a reduction proof can show that they are less favorable to the protocol.¹

Adversaries. We consider global passive *non-compromising* adversaries, that can observe all communication between protocol parties; and strictly stronger *partially compromising* (passive) adversaries, that can compromise protocol parties to learn the mapping between inputs and outputs for this party.

Anonymity Property. We leverage an indistinguishability based anonymity notion for sender anonymity: the adversary has to distinguish two senders of its own choosing [12], [13].

For a security parameter η , we say that a protocol achieves *strong anonymity*, if the adversary’s advantage remains negligible in η . Strong anonymity is relative to a strength η , which is bound to system parameters or analysis parameters such as the number of users or protocol parties, the latency overhead and the bandwidth overhead. These parameters typically increase as η increases, which improves the protocol’s anonymity.² Anonymity in relation to η unifies a wide variety of possible analyses on how the anonymity bound changes with changing system parameters, and user numbers and behaviors.

B. Brief Overview of the Proof Technique

As *non-compromising* adversaries are a subset of *partially compromising* adversaries, our proof technique for the former is a simplified case of the latter. In general, we derive our results in four main steps.

First, we define a concrete adversary \mathcal{A}_{paths} , that uses a well established strategy: upon recognizing the challenge message (as soon as it reaches a receiver) \mathcal{A}_{paths} constructs the possible paths this message could have taken through the network, and tries to identify the user who has sent the message.

Second, given the concrete adversary \mathcal{A}_{paths} , we identify a necessary invariant that any protocol has to fulfill in order to provide anonymity. Intuitively: *both challenge users chosen by the adversary must be active (i.e., send at least one message) before the challenge message reaches the recipient, and it must be possible for these messages to meet in at least one honest party along the way.* We prove that indeed this natural invariant is necessary for anonymity.

¹Such distributions might contain usage patterns, irregularities between users and synchronization failures that the adversary can exploit.

²In some analyses, individual parameters may reduce with increasing η , such as the bandwidth overhead per user, as the other parameters, such as the number of users, increase.

Next, we propose an ideal protocol Π_{ideal} that is optimal in terms of satisfying the invariant: The probability that Π_{ideal} fulfills the necessary invariant is at least as high as for any protocol within our model (limited by the same constraints for β and ℓ). Moreover, whenever Π_{ideal} satisfies the invariant, the advantage of \mathcal{A}_{paths} is zero. Thus, Π_{ideal} is at least as good as any protocol within our model at winning against \mathcal{A}_{paths} .

Finally, we calculate the advantage of \mathcal{A}_{paths} against Π_{ideal} to obtain a lower bound on the adversarial advantage against all protocols within our model.³

C. Scenarios and Lower Bounds

We devise necessary constraints for four different scenarios. Let Π be a protocol in our model, with N users, restricted by bandwidth overhead $\beta \in [0, 1]$ and latency overhead $\ell \geq 0$. For the *compromising* cases, the adversary can compromise c out of K protocol parties. We derive the following lower bounds for δ -sender anonymity in the respective scenarios.

Synchronized Users, Non-compromising Adversaries:

$$\delta \geq 1 - f_\beta(\ell), \text{ where } f_\beta(x) = \min\left(1, \left(\frac{x+\beta Nx}{N-1}\right)\right).$$

Synchronized Users, Partially Compromising Adversaries:

$$\delta \geq \begin{cases} 1 - [1 - (\frac{c}{\ell})/(\frac{K}{\ell})]f_\beta(\ell) & c \geq \ell \\ 1 - [1 - 1/(\frac{K}{c})]f_\beta(c) - f_\beta(\ell - c) & c < \ell. \end{cases}$$

Unsynchronized Users, Non-compromising Adversaries:

$$\delta \geq 1 - [1/2 + f_p(\ell)], \text{ where for } p \approx \beta \text{ we have } f_p(x) = \min(1/2, 1 - (1 - p)^x) \text{ for a positive integer } x.$$

Unsynchronized Users, Partially Compromising Adv.:

$$\delta \geq \begin{cases} 1 - [1 - (\frac{c}{\ell})/(\frac{K}{\ell})]^{[1/2 + f_p(\ell)]} & c \geq \ell \\ \left(1 - [1 - 1/(\frac{K}{c})]^{[1/2 + f_p(c)]}\right) \\ \times \left(1 - [1/2 + f_p(\ell - c)]\right) & c < \ell. \end{cases}$$

To keep the presentation concise, we focus on how to derive bounds for sender anonymity. As the bounds for recipient anonymity are obtained analogously, we only explain the adjustments in the proofs and the corresponding resulting bounds. The omitted canonical analysis can be found in [33].

D. Interpretation and Interesting Cases

Our first and third lower bounds, for respectively synchronized and unsynchronized user behaviors against in a non-compromised AC network, suggest an anonymity trilemma. Both lower bounds can be simplified under some natural constraints to the following simplified lemma:

Lemma 1 (Informal Trilemma). *For security parameter η , no protocol can achieve strong anonymity if $2\ell\beta < 1 - \epsilon(\eta)$, where $\epsilon(\eta) = \frac{1}{\eta^d}$ for any positive constant d .*

³ \mathcal{A}_{paths} is a possible adversary against all protocols within our model. If \mathcal{A}_{paths} has an advantage of δ against our ideal protocol Π_{ideal} (bounded by β and ℓ), then \mathcal{A}_{paths} will also have an advantage of at least δ against any protocol within our model (that is also bounded by β and ℓ). Thus, our bound for δ describes a lower bound on the adversarial advantage against any protocol within the model, while against particular protocols there can be other adversaries (in the same adversary class) with an even higher advantage.

Ideal asymptotic values for latency overhead is $\ell = O(1)$ (i.e., a constant number of hop separation from the receiver), while ideal asymptotic values for bandwidth overhead is $\beta = O(1/N) = O(1/poly(\eta))$ (i.e., a constant number of message per round from all $N = poly(\eta)$ users combined). It is easy to see that for this ideal overhead $\ell\beta = O(1/poly(\eta))$, the trilemma excludes strong anonymity, while, with latency overhead $\ell = N = O(poly(\eta))$ or with bandwidth overhead $\beta = O(1)$, the trilemma does not exclude strong anonymity.

We find some interesting possible overhead constraints for strong anonymity (e.g. $\ell = O(\eta)$ and $\beta = O(1/\eta)$) demanding some compromise in both latency and bandwidth. These constraints can help understand and improve existing AC protocols as well as inform the design of future AC protocols.

For partially compromised scenarios the requirements are naturally stronger. All constraints discussed for compromised case in the following part are in addition to the requirements from the non-compromised case. While bandwidth overhead might be sufficient against non-compromising adversaries, it is not sufficient if parts of the protocol are compromised. With $\ell = \eta$ and $\frac{K}{c} = constant$ strong anonymity may be possible, whereas with $\ell = O(1)$, strong anonymity is impossible, even for $K \in poly(\eta)$ and $c = O(1)$.

In case $c < \ell$, strong anonymity guarantees may be possible only if $2(\ell - c)p > 1 - \epsilon(\eta)$, where $p = p' + \beta$ combines the genuine user messages p' with their bandwidth overhead β . Our result shows a connection between the expected usage behavior p and the latency ℓ . If p is not particularly large, the latency cannot be low; otherwise, the path-length cannot be sufficiently high to ensure mixing at an honest node. In other words, unless p is very large (as should be the case for some file sharing applications), a low latency renders the AC protocol cheap to compromise, i.e., c can be low.

Our necessary constraints enable protocol designers of AC protocols to avoid bad trade-offs between latency and bandwidth overhead. For a given expected user behavior and a given target attacker against which the AC shall provide anonymity, our constraints clearly state which combinations of latency and bandwidth overhead to avoid.

E. Related Work

In contrast to previous work, our work provides necessary constraints for strong anonymity w.r.t. to bandwidth and latency overhead. While there is a successful line of work on provable anonymity guarantees [12], [26], [27], [34]–[37], it is incomparable since it provides lower bounds on anonymity for specific protocols, and does not prove any general statements about sufficient conditions for strong anonymity.

Previous work on attacks against anonymous communication protocols, except for Oya et al. [38], solely provides upper bounds on anonymity for specific protocols [39]–[42]. Oya et al. [38] cast their attack in a general model and provide a sophisticated generic attacker. However, they only compute bounds w.r.t. a dummy message rate against timed pool mixes, not against other protocols and not w.r.t. latency and compromisation rate. Even more important, none of these

results discuss the relationship of the lower bounds for latency and bandwidth overheads.

III. ANONYMITY DEFINITION AND USER DISTRIBUTIONS

A. AnoA-Style Anonymity Definition

We define our anonymity notions with a challenge-response game similar to AnoA [26], [27], where the challenger simulates the protocol and the adversary tries to deanonymize users. The challenger $\text{Ch}(\Pi, \alpha, b)$ allows the adversary to adaptively control user communication in the network, up to an uncertainty of one bit for challenges, and is parametric in the following parts: (i) the AC protocol Π to be analyzed, (ii) the so called *anonymity function* α , that describes the specific variant of anonymity such as sender anonymity, recipient anonymity and relationship anonymity, (iii) and the challenge bit b which determines the decision the challenger takes in challenge inputs from the adversary.

Given a security parameter η , we quantify the anonymity provided by the protocol Π simulated by $\text{Ch}(\Pi, \alpha, b)$ in terms of the advantage the probabilistic polynomial time (PPT) adversary \mathcal{A} has in correctly guessing Ch 's challenge bit b . We measure this advantage in terms of indistinguishability of random variables additively, where the random variables in question represent the output of the interactions $\langle \mathcal{A} | \text{Ch}(\Pi, \alpha, 0) \rangle$ and $\langle \mathcal{A} | \text{Ch}(\Pi, \alpha, 1) \rangle$.

Definition 1 ((α, δ) -IND-ANO). A protocol Π is (α, δ) -IND-ANO⁴ for the security parameter η , an adversary class \mathcal{C} , an anonymity function α and a distinguishing factor $\delta(\cdot) \geq 0$, if for all ppt machines $\mathcal{A} \in \mathcal{C}$, $\Pr[0 = \langle \mathcal{A} | \text{Ch}(\Pi, \alpha, 0) \rangle] \leq \Pr[0 = \langle \mathcal{A} | \text{Ch}(\Pi, \alpha, 1) \rangle] + \delta(\eta)$.

For an anonymity function α , we say that a protocol Π provides *strong anonymity* [12], [13] if it is (α, δ) -IND-ANO with $\delta \leq \text{neg}(\eta)$ for some negligible function neg . If δ is instead *non-negligible* in η , then we say that Π provides *weak anonymity*. Note that η does not measure the size of the anonymity set, but the computational limitation of the adversary.

Sender Anonymity. Sender anonymity characterizes the anonymity of users against a malicious server through the inability of the server (or some intermediary) to decide which of two *self-chosen* users have been communicating with the server. We borrow the sender anonymity α_{SA} definition from the AnoA framework [26], where α_{SA} selects one of two possible challenge users and makes sure that the users cannot be distinguished based on the chosen recipient(s) or message(s).

Definition 2 (Sender anonymity). A protocol Π provides δ -sender anonymity if it is (α_{SA}, δ) -IND-ANO for α_{SA} as defined in Figure 2.

Recipient Anonymity. Recipient anonymity characterizes that the recipient of a communication remains anonymous, even to observers that have knowledge about the sender in

⁴AnoA also allows a multiplicative factor ε ; we use the simplified version with $\varepsilon = 0$, such that δ directly corresponds to the adversarial advantage.

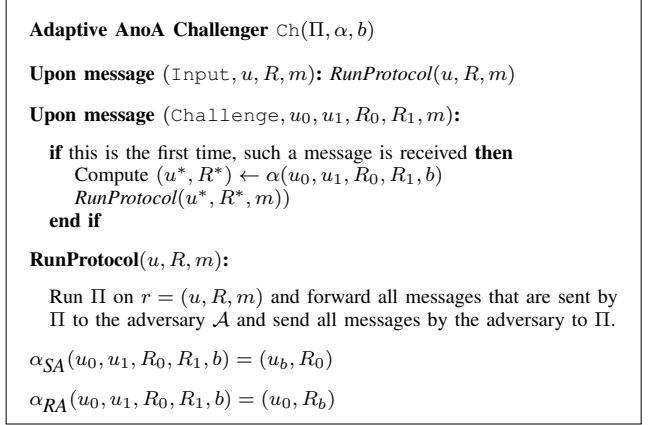


Fig. 2. Adaptive AnoA Challenger [26]

question. Similar to sender anonymity, we borrow the recipient anonymity α_{RA} definition from the AnoA framework, where α_{RA} selects one of two possible recipients for a message and makes sure that the recipients cannot be distinguished based on the chosen sender(s) or message(s).

Definition 3 (Recipient anonymity). A protocol Π provides δ -recipient anonymity if it is (α_{RA}, δ) -IND-ANO for α_{RA} as defined in Figure 2.

We omit the detailed technical notation of the anonymity functions in the following sections, and write $\Pr[0 = \mathcal{A}|b = i]$ instead of $\Pr[0 = \langle \mathcal{A} | \text{Ch}(\Pi, \alpha_{SA}, i) \rangle]$.

B. Game Setup

Let \mathcal{S} be the set of all senders, \mathcal{R} be the set of all recipients, and \mathcal{P} be the set of protocol parties that participate in the execution of the protocol (like relays/mix-nodes in Tor/mix-nets, for DC-net or P2P mixing users and protocol parties are the same). We consider a system of total $|\mathcal{S}| = N$ senders. Given our focus on *sender anonymity*, we need only a single element in \mathcal{R} . We allow the adversary to set the same entity (say R) as the recipient of all messages, and expect R to be compromised by the adversary. The adversary uses a challenge (as defined in Figure 2) of the form $(u_0, u_1, R, _, m_0)$, where $u_0, u_1 \in \mathcal{S}$, for our sender anonymity game.

We consider a completely connected topology, which means any party can send a message directly to any other party. We assume a standard (bounded) synchronous communication model as in [16], [17], [31], [32], where a protocol operates in a sequence of communication rounds.⁵ In each round, a party performs some local computation, sends messages (if any) to other party through an authenticated link. By the end of the round, every party receives all messages sent by the other parties to her the same round. With our focus on computing lower bounds, our model abstracts from the time

⁵While a time-sensitive model [43] would be more accurate, e.g., for low-latency protocols like Tor [44], such a model would only strengthen the attacker. As we present necessary constraints, our results also hold for the more accurate setting.

the computations at the node take and also the length of the messages. Nevertheless, as we are interested in quantifying the communication/bandwidth overhead, unlike [16], [17], [32], we do not assume that the parties have access to ready-made broadcast communication channels; Parties are expected to communicate with each other to implement broadcast features [31], [45]. Lastly, the use of the asynchronous communication model offers more capabilities to the attacker, and thus, our impossibility results for the synchronous model naturally apply to the asynchronous model as well.

We define the latency overhead ℓ as the number of rounds a message can be delayed by the protocol before being delivered. We define the bandwidth overhead β as the number of noise messages per user that the protocol can create in every round (i.e., the dummy message rate) and we do not restrict the time these noise messages reside within the protocol.

We consider two types of *global passive* adversaries: Our *non-compromising* adversaries (which model network-level eavesdroppers) can observe all communication between all protocol parties, but do not compromise any party of the AC protocol except the recipient R . We say that the AC protocol is *non-compromised*. Our strictly stronger *partially compromising* adversaries (which model hacking and infiltration capabilities) can additionally compromise some of the AC parties in the setup phase of the game to obtain these parties' mapping between the input messages and output messages during the protocol's runtime. We say that the AC protocol is *partially compromised*.

C. User Distributions

We consider two kinds of user distributions in our anonymity games and both of them assume an N sized set \mathcal{S} of users that want to send messages. In both cases, the adversary can choose any two senders $u_0, u_1 \in \mathcal{S}$. However, the time and method by which they actually send messages differs:

- In the *synchronized* user distribution the users globally synchronize who should send a message at which point in time. We assume that each user wants to send exactly one message. Consequently, we choose a random permutation of the set of users \mathcal{S} and the users send messages in their respective round. In every single round out of a total of N rounds exactly one user sends a message. Since the users globally synchronize their sending of messages, we allow the protocol to also globally decide on the bandwidth overhead it introduces. Note that here the requirements are identical to those of the Bulk protocol in [17].

- In the *unsynchronized* user distribution each of the N users wants to send messages eventually and we assume that each user locally flips a (biased) coin every round to decide whether or not to send a message. In this case we define the bandwidth overhead as an increased chance of users sending messages. Since the protocol does not globally synchronize the input messages, for noise messages also we allow the users to decide it locally and send noise messages with a certain probability.

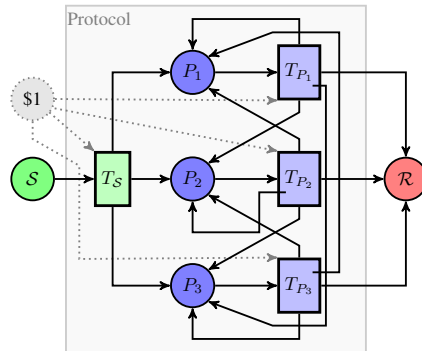


Fig. 3. Petri net of an AC protocol with $K = 3$ parties.

IV. A PROTOCOL MODEL FOR AC PROTOCOLS

An AC protocol allows any user in the set of users \mathcal{S} to send messages to any user in \mathcal{R} , via a set of anonymizing parties \mathcal{P} . We define protocols that are under observation of an eavesdropping adversary \mathcal{A} that may have compromised a set of c parties $\mathcal{P}_c \subseteq \mathcal{P}$ and that furthermore observes the communication links between any two parties, including users.

Technically, whenever a party $P_1 \in \mathcal{P} \cup \mathcal{S}$ sends a message to another party $P_2 \in \mathcal{P} \cup \mathcal{R}$, the adversary is able to observe this fact together with the current round number. However, we assume the protocol applies sufficient cryptography, s.t., the adversary can not read the content of any message except the messages sent to the malicious recipient, which technically results in simply being able to additionally recognize when the challenge reaches the recipient.

For an actual protocol, the sets \mathcal{S} , \mathcal{R} , and \mathcal{P} might not be mutually exclusive [15], [16], [18]. Since we have only one malicious party in \mathcal{R} , and the content of a message can only be read when it reaches its final recipient, we consider \mathcal{R} to be mutually exclusive from $\mathcal{S} \cup \mathcal{P}$ for the purpose of simplicity.

With the above preliminaries in mind, we shall now formally define our generic AC protocol using a petri net model.

A. Protocol Model

We model any AC protocol with K parties by a timed colored petri net [28]–[30] M , consisting of places \mathcal{S} for the users, P_1, \dots, P_K symbolizing the protocol parties, $\$1$ for randomness and R for recipients of messages, and colored tokens m symbolizing the messages (real or noise) sent by clients or protocol parties, and transitions T_S for inserting messages into the network and T_{P_1}, \dots, T_{P_K} as functions for sending the messages from one party to another. The structure of the petri net with its places, tokens and transitions remains the same for every AC protocol. However, the implementation of the guards within the transitions is different for different protocols: protocols can choose to which party messages are to be sent next and whether they should be delayed. But, protocols in M are oblivious to the challenge message or the challenge users. We refer to Figure 3 for a graphical depiction of petri net model M .

Definition 4 (Colored token). A colored token is represented by the tuple $m = \langle \text{msg}, \text{meta}, \text{t}_r, \text{ID}_t, \text{prev}, \text{next}, \text{ts} \rangle$, where, msg is the content of the message, meta is the internal protocol meta-data for this message, t_r is the time the message can remain in the network, ID_t is a new unique ID generated by each transition for each token by honest parties; dishonest parties instead keep ID_t untouched to allow the adversary to link incoming and outgoing messages, prev is party/user that sent the token and next is the user/party that receives the token. Finally, ts is the time remaining for the token to be eligible for a firing event (a feature of timed petri-net). Here, ts either describes when new messages are introduced into the petri net or is set to the next round, such that messages can be processed in every round as soon as they enter the network.

The four fields $\text{ID}_t, \text{prev}, \text{next}, \text{ts}$ are public, and are visible to the adversary. The remaining three fields msg, meta and t_r in a token are private and can not be observed by the adversary, with the exception that msg can be observed when a message reaches its destination, i.e, is received by a recipient. Formally, we introduce a set Tokens , that is initially empty and in which we collect the pair (t, r) , where t is a copy of a token and r the round number in which the token was observed.

Places. Any AC protocol with K parties $\mathcal{P} = \{P_1, \dots, P_K\}$ consists of the following places:

- \mathcal{S} : A token in \mathcal{S} denotes a user message (real or noise) which is scheduled to enter the network after ts rounds.
- $\mathcal{S}1$: This place is responsible for providing randomness. Whenever a transition picks a token from this place, the transition basically picks a random value.
- P_i with $P_i \in \mathcal{P}$: A token in P_i denotes a message which is currently held by the party $P_i \in \mathcal{P}$.
- R : A token in R denotes a message which has already been delivered to a recipient.

Transitions. As part of the *initial configuration*, the challenger populates \mathcal{S} on behalf of the protocol. All other places are initially empty. The transitions then consumes tokens from one place and generate tokens to other places, to modify the *configuration* of the petri-net. The event of consumption of a token from one place by a transition and generation of a new token represents the movement of a message from one party to another. We define the following transitions (refer to Figure 4 for the pseudocodes of the transitions):

- $T_{\mathcal{S}}$: takes a token $\langle \text{msg}, _, _, _, u, _, \text{ts} \rangle$ from \mathcal{S} and a token from $\mathcal{S}1$ to write $t = \langle \text{msg}, \text{meta}, \ell, \text{ID}_t, u, P_i, \text{ts} = 1 \rangle$ to P_i ; the values of i and meta are decided by the AC protocol.
- T_{P_i} : takes a token $\langle \text{msg}, \text{meta}, \text{t}_r, \text{ID}_t, _, P_i, \text{ts} \rangle$ from P_i and a token from $\mathcal{S}1$ to write $t = \langle \text{msg}, \text{meta}', \text{t}_r - 1, \text{ID}_t', P_i, P', 1 \rangle$ to P' . If P_i is an honest party ID_t' is freshly generated, but if P_i is a compromised party $\text{ID}_t' = \text{ID}_t$. The place $P' \in \{P_1, \dots, P_K\} \cup \{R\}$ and meta' are decided by the AC protocol, with the exception that if $\text{t}_r = 0$, P' always is R .

In either case, the transition also adds an element (t', r) to the set Tokens , where r is the current round number and t' is a copy of the respective (new) token t , with the fields meta and

$T_{\mathcal{S}}$ on tokens $q = \langle \text{msg}, _, _, _, u, _, \text{ts} \rangle$ from \mathcal{S} and $\mathcal{S}1$ from $\mathcal{S}1$:

$(P_i, \text{meta}) = f_{\Pi}(q, \mathcal{S}1)$; $\text{ID}_t =$ a fresh randomly generated ID
 $r =$ current round; $t = \langle \text{msg}, \text{meta}, \ell, \text{ID}_t, u, P_i, 1 \rangle$
if $P_i = R$ **then** $\text{Tokens} = \text{Tokens} \cup \langle \langle \text{msg}, _, _, \text{ID}_t, u, P_i, 1 \rangle, r \rangle$
else $\text{Tokens} = \text{Tokens} \cup \langle \langle _, _, _, \text{ID}_t, u, P_i, 1 \rangle, r \rangle$

Output: token t at P_i

T_{P_i} on tokens $q = \langle \text{msg}, _, \text{t}_r, \text{ID}_t, _, P_i, \text{ts} \rangle$ from P_i , $\mathcal{S}1$ from $\mathcal{S}1$:

$(P', \text{meta}') = f_{\Pi}(q, \mathcal{S}1)$; $r =$ current round
if $\text{t}_r - 1 = 0$ **then** $P' = R$
if P_i is honest **then** $\text{ID}_t' =$ a fresh randomly generated ID
else if P_i is compromised **then** $\text{ID}_t' = \text{ID}_t$
 $t = \langle \text{msg}, \text{meta}', \text{t}_r - 1, \text{ID}_t', P_i, P', 1 \rangle$
if $P_i = R$ **then** $\text{Tokens} = \text{Tokens} \cup \langle \langle \text{msg}, _, _, \text{ID}_t', P_i, P', 1 \rangle, r \rangle$
else $\text{Tokens} = \text{Tokens} \cup \langle \langle _, _, _, \text{ID}_t', P_i, P', 1 \rangle, r \rangle$

Output: token t at P'

f_{Π} : A function provided by Π to determine routing and the meta field.

Fig. 4. Transitions in petri net model M

t_r are removed. If the place where t was written to is not \mathcal{R} , then additionally the field msg is removed.

Game Setting. Recall that we define anonymity as a game between a PPT adversary \mathcal{A} and an honest challenger Ch .

Validity of the Protocol Model. The above protocol model M behaves as expected (more details in Lemma 2 in Appendix A). We show in Lemma 2 that the protocols indeed have a bandwidth overhead of β and a latency overhead of ℓ . For every message that is sent from one party in $\mathcal{S} \cup \mathcal{P}$ to another party in $\mathcal{P} \cup \mathcal{R}$, the adversary learns the time, the sender, and the receiver. When a message leaves the network, the attacker learns whether it was the target (i.e., the challenge) message. The attacker also learns the mapping between the input and output messages of compromised parties.

B. Expressing Protocols

Our protocol model M allows the expression of any AC protocol with very few, esoteric exceptions.

Mix networks can be naturally embedded into our model, in particular any stop-and-go mix [46] that uses discrete distribution and even AC protocols with specialized path selection algorithms [47], [48]. For the sake of our necessary constraints, low-latency protocols (with time-bounded channels) that are not round-based (e.g., Tor [44]) can be expressed in a round-based variant, since it only strengthens the protocols anonymity properties. This section illustrates embedding techniques into our model for some other kinds of protocols, but a much larger variety of protocols can be expressed in our model.

Users as protocol parties. In peer-to-peer protocols like dining cryptographers networks (DC net) [16], [18], there are no separate protocol parties, users act as a type of relays. Also, any noise sent by users counts into the bandwidth overhead of the protocol (we will see in Claim 2 that noise sent by nodes that are not users can be treated differently). Whenever a user wants to send a message it should use the transition $T_{\mathcal{S}}$, but when it acts as a relay it should use the transition T_{P_i} . For

interested readers, we show in Appendix A how to model a specific DC net type protocol using our petri net model.

Splitting and Recombining Messages. We model protocols that split and later re-combine messages by declaring one of the parts as the main message and the other parts as noise, which may count into the bandwidth overhead. This declaration is mainly required for the analysis, i.e., for evaluating the success of the adversary and for quantifying the amount of noise messages introduced by the protocol. We do not restrict the strategy by which the protocol decides which message is “the main share” (i.e., the message that is sent on) and which is “an additional share” (i.e., a fresh noise message). A more complex scenario involves threshold schemes in which a smaller number of shares suffices for reconstructing the message and in which some shares are dropped randomly. In such cases we consider the protocol to decide beforehand which of the constructed shares will be dropped later and to declare one of the remaining shares the “main share”.

Broadcasting Messages. If the protocol chooses to copy or broadcast messages to several receivers, we consider the copy sent to the challenge receiver to be the main message and copies sent to other receivers to be noise (which, if the copies are created by nodes that are not users, will not count into the bandwidth overhead).⁶

Private Information Retrieval. In schemes based on private information retrieval we require that the receiver retrieves the information sufficiently fast (within the latency limit). Otherwise, our method is similar to the broadcasting of messages: the receiver of interest will retrieve the main message, whereas other receivers will retrieve copies that are modeled as noise.

Excluded Protocols. For this work we exclude protocols that cannot guarantee the delivery of a message within the given latency bound (except if this occurs with a negligible probability). Moreover, we cannot easily express the exploitation of side channels to transfer information, e.g., sending information about one message in the meta-data of another message, or sending bits of information by not sending a message.

C. Construction of a Concrete Adversary

Given two challenge users u_0 and u_1 and the set of observed tokens $(t, r) \in \text{Tokens}$, where t is the token and r the round in which the token was observed, an adversary can construct the sets S_j (for $j \in \{0, 1\}$). Assume the challenge message arrives at the receiver R in a round r . We construct possible paths of varying length k , s.t., each element $p \in S_j$ represents a possible path of the challenge message starting from u_j ($j \in \{0, 1\}$) and the challenge message then arrives at R in round $r_k = r$. With challenge bit b , S_b cannot be empty, as the actual path taken by the challenge message to reach R has to be one element in S_b .

⁶We note that in some cases, where users act as nodes and broadcast messages to other users, our quantification of the bandwidth overhead might be a bit harsh. If the group of users to which the broadcast will be sent is known in advance (i.e., if messages are broadcast to all users or to pre-existing groups of users), we can allow the protocol to use a single receiver for these messages instead.

$$\begin{aligned}
S_j = \{ & p = (t_1.\text{prev}, \dots, t_k.\text{prev}, t_k.\text{next}) : \\
& ((t_1, r_1), \dots, (t_k, r_k)) \in \text{Tokens s.t.} \\
& t_1.\text{prev} = u_j \wedge t_k.\text{next} = R \\
& \wedge t_k.\text{msg} = \text{Challenge} \wedge k \leq \ell \\
& \wedge \forall_{i \in \{1, \dots, k-1\}} (t_i.\text{next} = t_{i+1}.\text{prev} \wedge r_{i+1} = r_i + 1 \\
& \wedge (\exists t'_{i+1} : (t'_{i+1}, r_{i+1}) \in \text{Tokens} \wedge t'_{i+1}.\text{prev} = t_i.\text{next} \\
& \wedge t'_{i+1}.\text{ID}_t = t_i.\text{ID}_t) \Rightarrow t'_{i+1} = t_{i+1}) \}
\end{aligned}$$

Definition 5 (Adversary $\mathcal{A}_{\text{paths}}$). *Given a set of users \mathcal{S} , a set of protocol parties \mathcal{P} of size K , and a number of possibly compromised nodes c , the adversary $\mathcal{A}_{\text{paths}}$ proceeds as follows: 1) $\mathcal{A}_{\text{paths}}$ selects and compromises c different parties from \mathcal{P} uniformly at random. 2) $\mathcal{A}_{\text{paths}}$ chooses two challenge users $u_0, u_1 \in \mathcal{S}$ uniformly at random. 3) $\mathcal{A}_{\text{paths}}$ makes observations and, based upon those, constructs the sets S_0 and S_1 . For any $i \in \{0, 1\}$, if $S_i = \emptyset$, then $\mathcal{A}_{\text{paths}}$ returns $1 - i$. Otherwise, it returns 0 or 1 uniformly at random.*

$\mathcal{A}_{\text{paths}}$ thus checks whether both challenge users *could* have sent the challenge message. We explicitly ignore differences in probabilities of the challenge users having sent the challenge message, as those probabilities can be protocol specific. Naturally, when $c = 0$, $\mathcal{A}_{\text{paths}}$ represents a *non-compromising* adversary; but when $c \neq 0$, $\mathcal{A}_{\text{paths}}$ is *partially compromising*.

D. Protocol Invariants

We now investigate the robustness of protocols against our adversary. We define an invariant that, if not satisfied, allows $\mathcal{A}_{\text{paths}}$ to win against any protocol. Moreover, we present a protocol that maximizes the probability of fulfilling the invariant. Moreover, we show that whenever the invariant is fulfilled by our protocol, the advantage of $\mathcal{A}_{\text{paths}}$ reduces to zero (as it is forced to randomly guess b).

Necessary invariant for protocol anonymity. It is necessary that at least both challenge users send messages in one of the ℓ rounds before the challenge message reaches the recipient, as otherwise there is no way both of them could have sent the challenge message. Moreover, on the path of the actual challenge message, there needs to be at least one honest (uncompromised) party, as otherwise the adversary can track the challenge message from the sender to the recipient (S_b will have exactly one element and S_{1-b} will be empty). Those two conditions together form our *necessary protocol invariant*.

Invariant 1. *Let u_0 and u_1 be the challenge users; let b be the challenge bit; and let t_0 be the time when u_b sends the challenge message. Assume that the challenge message reaches the recipient at r . Assume furthermore that u_{1-b} sends her messages (including noise messages) at $V = \{t_1, t_2, t_3, \dots, t_k\}$. Now, let $T = \{t : t \in V \wedge (r - \ell) \leq t < r\}$. Then,*

- (i) *the set T is not empty, and*
- (ii) *the challenge message passes through at least one honest node at some time t' such that, $t' \in \{\min(T), \dots, r - 1\}$.*

Claim 1 (Invariant 1 is necessary for anonymity). *Let Π be any protocol $\in \mathcal{M}$ with latency overhead ℓ and bandwidth*

overhead β . Let u_0, u_1, b and T be defined as in Invariant 1. If Invariant 1 is not satisfied by Π , then our adversary \mathcal{A}_{paths} as in Definition 5 wins.

We refer to Appendix B for the proof. We next claim that it suffices to consider noise messages sent by users that also remain within the system for at most ℓ rounds, i.e., noise messages that follow the same rules as real messages. Note that we consider every new message originating from any user's client as a fresh noise message.

Claim 2 (Internal noise does not influence Invariant 1). *Any message not originating from an end user $u \in \mathcal{S}$ does not influence the probability for Invariant 1 being true. Moreover, noise messages do not contribute to the probability for Invariant 1 being true after they stayed in the network for ℓ rounds.*

We refer to Appendix B for the proof. We henceforth consider noise messages as a protocol input.

E. Ideal Protocol

We construct a protocol Π_{ideal} that maximizes the probability of fulfilling Invariant 1. We show that the invariant is sufficient for Π_{ideal} to win against \mathcal{A}_{paths} , i.e., to reduce \mathcal{A}_{paths} 's advantage to 0. Claim 1 shows that for any protocol in our model \mathcal{A}_{paths} wins whenever Invariant 1 does not hold. Thus, an upper bound on the probability that Π_{ideal} satisfies Invariant 1 yields an upper bound for all these protocols.

Given the set of all protocol parties $P = \{P_0, \dots, P_{K-1}\}$ of size K , the strategy of Π_{ideal} is as follows: in a round r , Π_{ideal} delivers all messages scheduled for delivery to a recipient. All other messages (including the messages that enter Π_{ideal} in round r) are sent to the protocol party P_i with $i = r \bmod K$. For every message that enters the protocol, Π_{ideal} queries an oracle O for the number of rounds the message should remain in the protocol. We define the following events:

- $u.sent(x, y)$: user u has sent at least one message within rounds from x to y . For a single round we use $u.sent(x)$.
- $Cmpr(x)$: \mathcal{A}_{paths} has compromised the next x consecutive parties on the path.
- $\neg H$: NOT of event H .

Given a message sent at t_0 by sender x , and delivered to the recipient at $(t_0 + t)$, we define P_t for sender $v \in \mathcal{S} \setminus \{x\}$:

$$P_t = \sum_{j=r-\ell}^{t_0} \Pr[v.sent(j) \wedge \neg v.sent(j+1, t_0)] \times \Pr[\neg Cmpr(t)] \\ + \sum_{j=t_0+1}^r \Pr[v.sent(j) \wedge \neg v.sent(r-\ell, j-1)] \\ \times \Pr[\neg Cmpr(r-j)]$$

When $v = u_{1-b}$, and the message is the challenge message, P_t is the probability of fulfilling Invariant 1, for the strategy above. For each message, oracle O chooses an *optimal* t that maximizes the expectation of P_t over all users. After the oracle has decided the latencies for all messages, it sets the time t for the messages from u_{1-b} to ℓ . Since the oracle uses the knowledge of u_{1-b} , Π_{ideal} is slightly more powerful than protocols in M . Due to the over-approximation with this (not realizable) oracle, the resulting protocol is optimal w.r.t. Invariant 1 (Refer to Claim 3 and Claim 4).

Claim 3 (Ideal protocol is ideal for the invariant). *Against the given adversary \mathcal{A}_{paths} , Π_{ideal} satisfies Invariant 1 with probability at least as high as any other protocol in M .*

Claim 4 (Ideal protocol wins). *If Π_{ideal} satisfies Invariant 1, \mathcal{A}_{paths} has an advantage of zero:*

$$\Pr[b = \mathcal{A}_{paths} \mid \text{Invariant 1 holds}] = \frac{1}{2}$$

We refer to Appendix B for the proofs of Claim 3 and Claim 4.

V. SYNCHRONIZED USERS WITH NON-COMPROMISING ADVERSARIES

Our first scenario is a protocol-friendly user distribution U_B , where inputs from all users are globally synchronized: over the course of N rounds, exactly one user per round sends a message, following a random permutation that assigns one round to each user. Analogously, the protocol globally instructs the users to send up to $\beta \in [0, 1]$ noise messages **per user** per round, or $B = \beta N$ noise messages per round in total.

In real life, the user distribution is independent of the protocol. However, to make the user distribution protocol-friendly in our modeling we consider a globally controlled user distribution. For this scenario, we consider *non-compromising* passive adversaries that can observe all network traffic.

A. Lower Bound on Adversarial Advantage

Theorem 1. *For user distribution U_B , no protocol $\Pi \in M$ can provide δ -sender anonymity, for any $\delta < 1 - f_\beta(\ell)$, where $f_\beta(x) = \min(1, ((x + \beta Nx)/(N - 1)))$.*

Proof. By Claim 3 and Claim 4, we know that Π_{ideal} is an optimal protocol against \mathcal{A}_{paths} ; and with $c = 0$, \mathcal{A}_{paths} is our representative *non-compromising* adversary. Thus, it suffices to calculate the advantage of \mathcal{A}_{paths} against Π_{ideal} as a lower bound of the adversary's advantage against any protocol.

Let, u_0 and u_1 be the users chosen by the adversary and let b be the challenge bit. Let t_0 be the round in which u_b sends the challenge message and let r be the round in which the challenge message reaches the recipient.

Recall that Invariant 1 is necessary for the protocol to provide anonymity; u_{1-b} sends her messages (can be a noise message) at $V = \{t_1, t_2, t_3, \dots, t_k\}$, then $T = \{t : t \in V \wedge (r - \ell) \leq t < r\}$. Since we are considering a non-compromising adversary, $\Pr[\text{Invariant 1 is true}] = \Pr[T \text{ is not empty}]$.

With the above in mind, let us define the following events:

- H_1 : In ℓ rounds u_{1-b} sends at least one noise message.
- H_2 : u_{1-b} sends his own message within the chosen ℓ rounds.
- H_3 : there is at least one message from u_{1-b} within the chosen ℓ rounds $\equiv T$ is not empty \equiv Invariant 1 is true.

Consider any slice of ℓ rounds around the challenge message, there are exactly $(\ell - 1)$ user messages other than the challenge message. Hence, any slice of ℓ rounds yields the same probability of containing a user message from u_{1-b} , except when $r < \ell$ OR $r > N$ where the probability is smaller. Thus, no matter what value of t is returned by O , $\Pr[H_2] \leq \frac{\ell-1}{N-1}$.

Given any values $\ell, \beta \geq 0$, \mathcal{A}_{paths} has the least chance of winning, if for a given interval of ℓ rounds, $\beta N \ell$ unique users

are picked to send the noise messages in such a way that they are not scheduled to send their own messages in that interval.

$$\Pr[\neg H_3] = \Pr[\neg H_1, \neg H_2] \geq \max(0, (N - \ell - \beta N \ell) / (N - 1)).$$

$$\Pr[H_3] = 1 - \Pr[\neg H_3] \leq \min(1, ((\ell + \beta N \ell) / (N - 1))).$$

Thus, we can bound the probability for the adversary as $\Pr[0 = \mathcal{A}_{paths}|b = 1] = \Pr[1 = \mathcal{A}_{paths}|b = 0] = \frac{1}{2}\Pr[H_3]$; and $\Pr[0 = \mathcal{A}_{paths}|b = 0] = 1 - \frac{1}{2}\Pr[H_3]$. And therefore, since $\delta \geq \Pr[0 = \mathcal{A}_{paths}|b = 0] - \Pr[0 = \mathcal{A}_{paths}|b = 1]$, $\delta \geq 1 - \Pr[H_3] \geq 1 - f_\beta(\ell)$. \square

B. Impossibility for Strong Anonymity

We now investigate under which constraints for ℓ and β Theorem 1 rules out strong anonymity.

Theorem 2. For user distribution U_B with $\ell < N$ and $\beta N \geq 1$, no protocol $\Pi \in M$ can achieve strong anonymity if $2\ell\beta < 1 - \epsilon(\eta)$, where $\epsilon(\eta) = \frac{1}{\eta^2}$ for a positive constant d .

We refer to Appendix B for the proof.

Interesting Cases. For illustration, we now discuss a few examples for different values of ℓ , β , and N .

1) If $\ell = N$, we can have $\delta = 0$ even for $\beta = 0$. Anonymity can be achieved trivially by accumulating all messages from all N users and delivering them together at round $(N + 1)$. In this case $2\ell\beta = 0 < 1 - \epsilon(\eta)$, but also $\beta N = 0 < 1$.

2) $\beta = \frac{1}{\eta}$, $\ell = \eta$: We have $\delta \geq \frac{N - \eta - N}{N} \geq \frac{-\eta}{N}$. In ℓ rounds the protocol can send $\ell\beta N = N$ noise messages and achieve strong anonymity (all N users send a noise message each).

3) $\beta = \frac{1}{2\tau}$, $\ell = \tau$, where τ is a positive integer: Here we have, $\delta \geq \frac{N - \tau - N}{N} = \frac{1}{2} - \frac{\tau}{N}$. Here, strong anonymity is possible if $\frac{\tau}{N} \geq \frac{1}{2} - \text{neg}(\eta)$. Even though $2\ell\beta = 1 > 1 - \text{neg}(\eta)$, anonymity depends on the relation between τ and N .

4) $\beta = \frac{1}{9}$, $\ell = 3$: For $\eta > 3$ and $N > 4$, which is a very natural assumption, we have $2\ell\beta = \frac{2}{3} < 1 - \text{neg}(\eta)$. Then, $\delta \geq \frac{N - 3 - \frac{N}{3}}{N} > \text{neg}(\eta)$. In ℓ rounds Π_{ideal} receives only $(\frac{N}{3} + 3)$ messages, and thus, with high probability u_{1-b} does not send a message. Hence, Π_{ideal} cannot achieve strong anonymity.

VI. SYNCHRONIZED USERS WITH PARTIALLY COMPROMISING ADVERSARIES

We now extend our analysis of the previous section by having compromised protocol parties. Given the set of protocol parties P , now our adversary \mathcal{A}_{paths} can compromise a set of c parties $P_c \subset P$. If \mathcal{A}_{paths} can compromise all the parties in P , anonymity is broken trivially - that's why we do not analyze that case separately. Recall from Section IV-C that \mathcal{A}_{paths} picks the c parties from P uniformly at random. We consider the same user distribution U_B as in Section V.

A. Lower Bound on Adversarial Advantage

Theorem 3. For user distribution U_B , no protocol $\Pi \in M$ can provide δ -sender anonymity, for any

$$\delta < \begin{cases} 1 - [1 - \binom{c}{\ell} / \binom{K}{\ell}] f_\beta(\ell) & c \geq \ell \\ 1 - [1 - 1 / \binom{K}{c}] f_\beta(c) - f_\beta(\ell - c) & c < \ell \end{cases}$$

where $f_\beta(x) = \min(1, ((x + \beta N x) / (N - 1)))$.

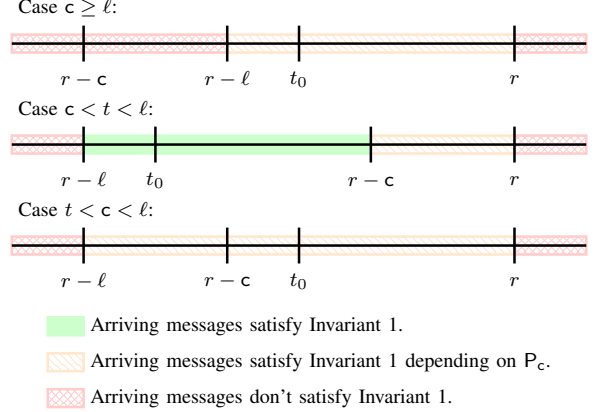


Fig. 5. Satisfying Invariant 1 depending on the arrival time of messages from u_{1-b} in the cases of the proof for Theorem 3.

Proof. Let u_0, u_1 be the challenge users and let b be the challenge bit. Moreover, let t_0 be the time the challenge message is sent by u_b and let $r = t_0 + t$ be the time it is received by the recipient, where t is the delivery time decided by the oracle O . Similar to Section V, we now calculate the advantage of \mathcal{A}_{paths} against Π_{ideal} .

We distinguish two cases, depending on ℓ and c : 1) First, where the number of compromised parties c is at least as large as the maximal latency ℓ . In this case, all parties on the path of the challenge message could be compromised. 2) Second, where not all parties on the path of the challenge message can be compromised. And hence, the analysis focuses on the arrival times of messages from u_{1-b} . For a graphical depiction of the relationship between the rounds a message from u_{1-b} arrives and it satisfying Invariant 1 we refer to Figure 5.

1) Case $c \geq \ell$. We know, $\ell \geq t$ holds by definition. The invariant is true only if u_{1-b} sends at least one message in one of the rounds between $(r - \ell)$ and $(r - 1)$. Additionally, if u_{1-b} sends at least one message in $\{r - \ell, \dots, t_0\}$, the invariant holds only if there is at least one non-compromised party on the path between t_0 and $(r - 1)$. Whereas, if u_{1-b} does not send any message in $\{r - \ell, \dots, t_0\}$, and the first message from u_{1-b} in the interval $\{t_0 + 1, r - 1\}$ arrives at t_1 , the invariant holds only if there is at least one non-compromised party on the path between t_1 and $(r - 1)$.

Note that $K > c \geq \ell$. Also recall from Section IV that \mathcal{A}_{paths} picks the c parties uniformly at random from K parties. Hence,

$$\begin{aligned} & \Pr[\text{Invariant 1 is true}] \\ & \leq \sum_{j=r-\ell}^{t_0} \Pr[u_{1-b}.\text{sent}(j) \wedge \neg u_{1-b}.\text{sent}(j+1, t_0)] \\ & \quad \times \Pr[\neg \text{Cmpr}(t)] \\ & + \sum_{j=t_0+1}^r \Pr[u_{1-b}.\text{sent}(j) \wedge \neg u_{1-b}.\text{sent}(r-\ell, j-1)] \\ & \quad \times \Pr[\neg \text{Cmpr}(r-j)] \\ & \leq \Pr[\neg \text{Cmpr}(\ell)] \times \Pr[u_{1-b}.\text{sent}(r-\ell, r-1)] \\ & \leq [1 - \binom{c}{\ell} / \binom{K}{\ell}] \times \min(1, ((\ell + \beta N \ell) / (N - 1))). \end{aligned}$$

By Claim 1 the adversary wins whenever Invariant 1 is not true, and by Claim 4 \mathcal{A}_{paths} has zero advantage whenever Π_{ideal} satisfies the invariant. Hence, we know that the probability that the adversary guesses incorrectly is bounded by:

$$\Pr[0 = \mathcal{A}_{paths}|b = 1] = \Pr[1 = \mathcal{A}_{paths}|b = 0] \\ \leq \frac{1}{2} \Pr[\text{Invariant 1 is true}] \leq \frac{1}{2} [1 - \binom{c}{\ell} / \binom{K}{\ell}] \times \min(1, (\frac{\ell + \beta N \ell}{N-1})).$$

Thus, $\delta \geq 1 - [1 - \binom{c}{\ell} / \binom{K}{\ell}] \times \min(1, (\frac{\ell + \beta N \ell}{N-1}))$.

2) Case $c \leq \ell$: The probability that all parties on the mutual path of the challenge message and a message from the alternative sender u_{1-b} are compromised now mainly depends on the arrival time of the messages from u_{1-b} . We find two sub-cases depending on the oracle's choice for t .

2a) Case $c \leq t$:

$$\Pr[\text{Invariant 1 is true}] \\ \leq \Pr[u_{1-b}.\text{sent}(r - \ell, r - c)] + \Pr[\neg u_{1-b}.\text{sent}(r - \ell, r - c)] \\ \times \Pr[u_{1-b}.\text{sent}(r - c, r)] \times \Pr[\neg \text{Cmpr}(c)] \\ \leq \min(1, (\frac{(\ell - c) + \beta N(\ell - c)}{N-1})) \\ + \min(1, (\frac{N - (\ell - c) - \beta N(\ell - c)}{N-1})) (\frac{c + \beta N c}{N - (\ell - c) - \beta N(\ell - c)}) [1 - \frac{1}{\binom{K}{c}}] \\ \leq f_\beta(\ell - c) + f_\beta(c) [1 - 1/\binom{K}{c}].$$

Note that the probability that there are no messages from u_{1-b} in $[(r - \ell), (r - c)]$ and that there is at least one message from u_{1-b} in $[(r - c), r]$ are not independent from each other. The best thing a protocol can do with the noise messages is to have $N\beta\ell$ unique users, different from the ℓ users who send their actual message, send the noise messages. Thus, if a user sends a message in $[(r - \ell), (r - c)]$, he can not send a message in $[(r - c), r]$. The above calculations are done considering that best scenario. Also note that the value of K may be larger or smaller than ℓ and t , but as long as $c \leq K$, the bound given above holds. Hence, $\delta \geq 1 - f_\beta(\ell - c) - [1 - 1/\binom{K}{c}] \times f_\beta(c)$.

2b) Case $t < c$:

$$\Pr[\text{Invariant 1 is true}] \\ \leq \Pr[u_{1-b}.\text{sent}(r - \ell, r - c)] \times \Pr[\neg \text{Cmpr}(t)] \\ + \Pr[\neg u_{1-b}.\text{sent}(r - \ell, r - c)] \\ \times \Pr[u_{1-b}.\text{sent}(r - c, r)] \times \Pr[\neg \text{Cmpr}(t)] \\ \leq \Pr[u_{1-b}.\text{sent}(r - \ell, r - c)] + \Pr[\neg u_{1-b}.\text{sent}(r - \ell, r - c)] \\ \times \Pr[u_{1-b}.\text{sent}(r - c, r)] \times \Pr[\neg \text{Cmpr}(t)]$$

The event expression above is the same as in the previous case ($t > c$). The bound on δ thus follows analogously. \square

B. Impossibility for Strong Anonymity

Theorem 4. For user distribution U_B with $K \in \text{poly}(\eta)$, $K > c \geq \ell$, $\ell < N$ and $\beta N \geq 1$, no protocol $\Pi \in M$ can achieve strong anonymity if $2\ell\beta < 1 - \epsilon(\eta)$ or $\ell \in \mathcal{O}(1)$, where $\epsilon(\eta) = 1/\eta^d$ for a positive constant d .

We refer to Appendix B for the proof. To achieve strong anonymity against \mathcal{A}_{paths} , we need $\ell \in \omega(1)$, additional to the constraint of $2\ell\beta > 1 - \text{neg}(\eta)$. We now focus on the constraint $\ell \in \omega(1)$ and refer to Section V-B for a comprehensive case study on the other constraint.

Interesting Cases. Now we are going to discuss a few interesting cases for different values of $\ell < c$, and K .

1) $\ell = \eta$ and $K/c = \text{constant}$: In this case we have, $\binom{c}{\ell} / \binom{K}{\ell} = \frac{c(c-1)\dots(c-\ell+1)}{K(K-1)\dots(K-\ell+1)} < (c/K)^\ell = (c/K)^\eta$. Hence, $\binom{c}{\ell} / \binom{K}{\ell}$ becomes negligible and strong anonymity is possible. Even though c has a high value, because of the high value of ℓ it is highly likely that the challenge message will meet a message from u_{1-b} at some honest node, given a high value of β such that $2\ell\beta > 1 - \text{neg}(\eta)$.

2) $\ell = \mathcal{O}(1)$, $c = \mathcal{O}(1)$: Now we have, $\binom{c}{\ell} / \binom{K}{\ell} = \frac{c(c-1)\dots(c-\ell+1)}{K(K-1)\dots(K-\ell+1)} > ((c-\ell)/(K-\ell))^\ell$. But $K \in \text{poly}(\eta)$, and c and ℓ can only have integer values. Hence $((c-\ell)/(K-\ell))^\ell$ is non-negligible, and hence $\binom{c}{\ell} / \binom{K}{\ell}$ is also non-negligible. Even though c has a small value, ℓ is also small. Hence, it is unlikely that the challenge message will mix with a message from u_{1-b} at some honest node. Thus, strong anonymity cannot be achieved.

Theorem 5. For user distribution U_B with $K \in \text{poly}(\eta)$, $c \in \mathcal{O}(1)$, $K > \ell > c$, $\ell < N$ and $\beta N \geq 1$, no protocol $\Pi \in M$ can achieve strong anonymity if $2(\ell - c)\beta < 1 - \epsilon(\eta)$, where $\epsilon(\eta) = \frac{1}{\eta^d}$ for a positive constant d .

We refer to Appendix B for the proof. The analysis in this case is exactly same as Section V-B, except that here we need to consider the slice of $(\ell - c)$ rounds instead of ℓ rounds.

It is worth repeating here, all the constraints we have derived in Section V and Section VI are necessary for anonymity, but they are not sufficient conditions for anonymity.

VII. UNSYNCHRONIZED USERS WITH NON-COMPROMISING ADVERSARIES

In this and the subsequent section we use an unsynchronised user distribution U_P : In each round, independent of other users and other rounds, each client tosses a biased coin with success probability $p \in (0, 1]$. On a success the client sends a message in that round, otherwise it does not send a message. Consequently, the number of messages per round follows Binomial distribution $\text{Binom}(N, p)$ if the number of users N is large and p sufficiently small, the resulting binomial distribution reduces to a Poisson distribution, which is a close approximation of real-life traffic patterns.

For a protocol with bandwidth overhead β , we distinguish between the actual probability that users want to send messages p' and the value for p that we use in our analysis, i.e., we set $p = p' + \beta$. In this unsynchronised scenario the bandwidth of genuine messages contributes to the anonymity bound. As in Section V we consider a *non-compromising* adversary.

A. Lower Bound on Adversarial Advantage

Theorem 6. For user distribution U_P , no protocol $\Pi \in M$ can provide δ -sender anonymity, for any $\delta < 1 - (\frac{1}{2} + f_p(\ell))$, where $f_p(x) = \min(1/2, 1 - (1 - p)^x)$ for a positive integer x .

Proof. Since we consider a non-compromising adversary, $\Pr[\text{Invariant 1 is True}] = \Pr[T \text{ is not empty}]$, where T is defined as in Invariant 1.

Let us consider the random variables $X^{(1)}, X^{(2)}, \dots, X^{(N)}$, where $X^{(i)}$ denotes the event of the i^{th} user sending her own

message within a given interval of ℓ rounds $[a, b]$, with $(b - a) = \ell$. All $X^{(i)}$ s are mutually independent and we have,

$$X^{(i)} = \begin{cases} 0 & \text{with probability } (1-p)^\ell \\ 1 & \text{with probability } (1-(1-p)^\ell). \end{cases}$$

Next, let $X = \sum_{i=1}^N X^{(i)}$ be a random variable representing the number of users that send messages in an interval of ℓ rounds. We calculate for the expected value $\mathbb{E}[X]$ of X ,

$$\mathbb{E}[X] = \sum_{i=1}^N \mathbb{E}[X^{(i)}] = N(1 - (1-p)^\ell) = \mu.$$

Using the Chernoff Bound on the random variable X we derive $\Pr[X - \mu \geq Na] \leq \exp(-2a^2N)$, which for $a = \frac{\mu}{N}$ lets us estimate, $\Pr[X \geq 2\mu] \leq \exp(-2(\mu^2/N^2)N)$. For brevity in the following calculation we denote, $\Pr[X \geq 2\mu]$ by E and the event that T is non-empty by Y and since all users are acting independently from each other we get for $j \in \{0, \dots, N\}$, $\Pr[Y|X = j] = 1 - \Pr[-Y|X = j] = \frac{j}{N}$.

For $2\mu \leq N$, we have,

$$\begin{aligned} & \Pr[Y] \\ &= \Pr[X \geq 2\mu] \times \Pr[Y|X \geq 2\mu] + \Pr[X < 2\mu] \times \Pr[Y|X < 2\mu] \\ &\leq \Pr[X \geq 2\mu] \times \Pr[Y|X = N] + \Pr[X < 2\mu] \times \Pr[Y|X = 2\mu] \\ &= E \times \Pr[Y|X = N] + (1-E) \times \Pr[Y|X = 2\mu] \\ &= E \times \frac{N}{N} + (1-E) \times \frac{2\mu}{N} = 1 - (1-E)(1-2f_p(\ell)). \end{aligned}$$

If $2\mu > N$, we get with $f(\ell) = \min(\frac{1}{2}, 1 - (1-p)^\ell)$, $\Pr[Y] \leq E + (1-E)1 \leq 1 \leq 1 - (1-E)(1-2f_p(\ell))$.

Thus, $\delta \geq 1 - \Pr[Y] \geq (1-E)(1-2f_p(\ell))$. We now use Markov's Inequality on X and derive $E = \Pr[X \geq 2\mu] \leq \frac{1}{2}$, which means, $\delta \geq \frac{1}{2}(1-2f_p(\ell)) \geq \frac{1}{2} - f_p(\ell)$. \square

Note that in the proof of Theorem 6, in case p is a constant and N is a very high value, then E goes towards zero and instead of using Markov's inequality, we can derive $\delta \geq 1 - 2f_p(\ell)$.

B. Impossibility for Strong Anonymity

Theorem 7. For user distribution U_P and $p > 0$, no protocol $\Pi \in M$ can achieve strong anonymity if $2\ell p < 1 - \epsilon(\eta)$, where $\epsilon(\eta) = 1/\eta^d$ for a positive constant d .

We refer to Appendix B for the proof. Similar to the constraints in Section V and Section VI, this is also a necessary constraint for anonymity, not a sufficient condition. There can exist ℓ and p such that $2\ell p > 1 - \text{neg}(\eta)$, but still no protocol can achieve strong anonymity.

Interesting Cases. Now we are going to discuss a few interesting cases for different values of ℓ , p , and N .

1) $p = \frac{1}{\eta}$, $\ell = \eta$: Here, $f_p(\ell) = 1 - (1-p)^\ell > 1 - 1/e > \frac{1}{2}$. Hence, $\delta \geq \frac{1}{2} - f_p(\ell) = 0$. Since $p\ell = 1$, in ℓ rounds the protocol has 1 message per user on an average. So, the protocol has a high chance of winning. Whereas in Section V-B, we saw that Π_{ideal} can win with absolute certainty in this case.

2) $p = \frac{1}{2\tau}$, $\ell = \tau$, τ is a positive integer: even for $\tau > 2$, $f_p(\ell) = 1 - (1-p)^\ell < 0.45$. Hence, $\delta \geq \frac{1}{2} - f_p(\ell) > 0.05$. Even though $2\ell p = 1$, strong anonymity can not be achieved. In an expected scenario, in a slice of ℓ rounds only $p\ell = \frac{1}{2}$ portion of the total users send messages, and hence there is a significant chance that u_{1-b} is in the other half. Note that this is different

from the scenario with synchronized users where Π_{ideal} could achieve strong anonymity in this case (c.f. Section V-B).

3) $p = \frac{1}{9}$, $\ell = 3$: Here, $f_p(\ell) = 1 - (1-p)^\ell = 1 - (\frac{8}{9})^3 < 0.29$, and $\delta \geq \frac{1}{2} - f_p(\ell) > 0.21$; because of low values of both p and ℓ only a few users send messages within the interval of ℓ rounds, and hence the protocol has a small chance to win. As in Section V-B, Π_{ideal} can not achieve strong anonymity in this case, since the necessary constraints are not satisfied.

VIII. UNSYNCHRONIZED USERS WITH PARTIALLY COMPROMISING ADVERSARIES

Finally, we consider partially compromising adversaries that can compromise a set of c parties $P_c \subset P$ for the user distribution U_P defined in Section VII.

A. Lower Bound on Adversarial Advantage

Theorem 8. For user distribution U_P , no protocol $\Pi \in M$ can provide δ -sender anonymity, for any

$$\delta < \begin{cases} 1 - [1 - \binom{c}{\ell} / \binom{K}{\ell}] [\frac{1}{2} + f_p(\ell)] & c \geq \ell \\ \left(1 - [1 - 1/\binom{K}{c}] [\frac{1}{2} + f_p(c)] \right) \\ \times \left(1 - [1/2 + f_p(\ell - c)] \right) & c < \ell \end{cases}$$

where $f_p(x) = \min(1/2, 1 - (1-p)^x)$ for a positive integer x .

We derive the bound in Theorem 8 by combining the techniques presented in Section VI and Section VII. Since the proof does not introduce novel techniques, we omit it and instead refer the interested reader to Appendix B for the proof.

B. Impossibility for Strong Anonymity

To analyze the negligibility condition of δ in this scenario, we heavily borrow the analyses that we already have conducted in Section VII-B and Section VI-B. We are going to analyze this scenario in two parts:

Case $c \geq \ell$: We have, $\delta \geq 1 - [1 - \binom{c}{\ell} / \binom{K}{\ell}] [\frac{1}{2} + f_p(\ell)]$.

To make δ negligible, both the factors $[1 - \binom{c}{\ell} / \binom{K}{\ell}]$ and $[1/2 + f_p(\ell)]$ have to become overwhelming. From Theorem 4, we know that we need $\ell \in \omega(1)$ to make $[1 - \binom{c}{\ell} / \binom{K}{\ell}]$ overwhelming. This is a necessary condition, but not sufficient. For a detailed discussion, we refer to Section VI-B. From Section VII-B we know that the necessary condition for $[1/2 + f_p(\ell)]$ to be overwhelming is $2\ell p > 1 - \text{neg}(\eta)$. Hence, both conditions are necessary to achieve strong anonymity.

Case $c < \ell$: We have,

$$\delta \geq (1 - [1/2 + f_p(\ell - c)])(1 - [1 - 1/\binom{K}{c}] [\frac{1}{2} + f_p(c)]).$$

In the above expression, we can see two factors:

(i) $F_1 = (1 - [1/2 + f_p(\ell - c)])$, (ii) $F_2 = (1 - [1 - 1/\binom{K}{c}] [\frac{1}{2} + f_p(c)])$.

To make δ negligible, it suffices that F_1 or F_2 become negligible. Unlike Section VI, here $f_p(\ell - c)$ and $f_p(c)$ are independent, which allows us to analyze F_1 and F_2 independently. First, F_1 is similar to the δ -bound in Section VII, except that we consider $f_p(\ell - c)$ instead of $f_p(\ell)$. Hence, the analysis of F_1 is analogous to Section VII-B. Second, F_2 is negligible if both $[1 - 1/\binom{K}{c}]$ and $[1/2 + f_p(c)]$ are overwhelming. From Section VI-B we know that $[1 - 1/\binom{K}{c}]$ can not be overwhelming for a constant c . Moreover, $f_p(c)$ can be analyzed exactly as $f_p(\ell)$ in Section VII-B.

IX. RECIPIENT ANONYMITY

We derive impossibility results for recipient anonymity analogous to our results for sender anonymity via the same strategy we employed in the previous sections. In this case, since we are considering recipient anonymity, we assume only one sender in \mathcal{S} , and N' users in \mathcal{R} . Here, the adversary is naturally not informed about the delivery of the challenge message by a recipient, but of the sending of the challenge message by the sender. Moreover, instead of ignoring all internally generated messages in Claim 2 we ignore all internally terminating messages. Note that this gives β a slightly different flavor.

Synchronized Users. We slightly tweak the user distribution to suit the definition of *recipient anonymity*. We assume that all the input messages come within N' rounds, exactly one message per round, following a random permutation that assigns one round to each recipient. In a given round, the sender sends a message to the assigned recipient. Then, the protocol decides when to deliver the message to the recipient, but not delaying more than ℓ rounds. Let $f_\beta^{RA}(x) = \min\left(1, \left(\frac{(x+\ell)+(x+\ell)\beta N'}{N'}\right)\right)$. Then we get that no protocol $\Pi \in M$ can provide δ -recipient anonymity in the following cases:

- Without compromisation: $\delta < 1 - f_\beta^{RA}(\ell)$.
- For adversaries that compromise up to c parties:
 - if $c \geq \ell$: $\delta < 1 - [1 - \binom{c}{\ell} / \binom{K}{\ell}] f_\beta^{RA}(\ell)$.
 - if $c < \ell$: $\delta < 1 - [1 - 1/\binom{K}{c}] f_\beta^{RA}(c) - f_\beta^{RA}(\ell - c)$.

Moreover, no protocol M with $K \in \text{poly}(\eta)$ can achieve strong recipient anonymity when $\ell < N'$ and $\beta N' \geq 1$ in the following cases, where $\epsilon(\eta)$ is a non-negligible function.

- Without compromisation: if $4\ell\beta < 1 - \epsilon(\eta)$,
- For adversaries that compromise up to c parties:
 - if $K > c \geq \ell$: $4\ell\beta < 1 - \epsilon(\eta)$ OR $\ell \in \mathcal{O}(1)$.
 - if $K > \ell > c$: $4(\ell - c)\beta < 1 - \epsilon(\eta)$.

Unsynchronized Users. Similar to the previous case, here also we borrow the definition of user distribution from Section VII, with minor modifications. The biased coins are now associated with recipients instead of senders — in each round the sender sends a message **for a recipient**, with probability p . Let $f_p^{RA}(x) = \min(1/2, 1 - (1-p)^{\ell+x})$. Then we get that no protocol $\Pi \in M$ can provide δ -recipient anonymity in the following cases:

- Without compromisation: $\delta < 1 - (1/2 + f_p^{RA}(\ell))$.
- For adversaries that compromise up to c parties:
 - If $c \geq \ell$: $\delta < [1 - \binom{c}{\ell} / \binom{K}{\ell}] [1/2 + f_p^{RA}(\ell)]$.
 - If $c < \ell$: $\delta < \left(1 - [1/2 + f_p^{RA}(\ell - c)]\right) \times \left(1 - [1/2 + f_p^{RA}(c)] [1 - 1/\binom{K}{c}]\right)$.

Moreover, for $p > 0$, no protocol can achieve strong recipient anonymity if $2\ell p < 1 - \epsilon(\eta)$, where $\epsilon(\eta)$ is a non-negligible function. For a detailed recipient-anonymity analysis, we refer the readers to the extended version [33].

X. IMPLICATIONS

To put our result into perspective, we discuss whether our trilemma excludes strong anonymity for a few AC protocols

from the literature. More precisely, this section exemplarily applies the results from Theorem 2 and Theorem 7, i.e., with synchronized and unsynchronized user distributions and a global network-level, non-compromising adversary. We use both results since for some AC protocols (e.g., DC-nets [15]) the synchronized user distribution is more accurate and for other protocols (e.g., Tor [10]) the unsynchronized user distribution is more accurate. Our constraints mark an area on a 2D graph (see Figure 6) with latency overhead (x-axis) versus bandwidth overhead (y-axis) where strong anonymity is impossible. As the latency of some AC protocols depends on system parameters and we want to place the protocols in a 2D graph, we carefully choose system parameters and make a few simplifying assumptions, which are subsequently described.

This section is solely intended to put our impossibility result into perspective. It is not meant and not qualified to be a performance and scalability comparison of the discussed AC protocols. Table I in the appendix summarizes bounds on the bandwidth β and latency overhead ℓ (in the sense of this work).

Technically, this section considers translations of AC protocols into our protocol model. As these translations do not provide any additional insights, we do not present the full translated protocols but only the abstraction steps. We abstract away the cryptographic instantiation of messages including the bandwidth overhead they introduce over the plaintext. We assume an upper bound on the latency of the protocol and are oblivious to server-side noise (see Claim 2). Moreover, recall that we are only interested in the question whether our trilemma excludes strong anonymity for the ten AC protocols from the literature; hence, we consider the upper bound on the latency and bandwidth overhead for deterministic latency. For randomized latency, such as Loopix [24], we list for simplicity the expected delay as the latency bound.

Low-latency protocols such as Tor [10], Hornet [49], and Herd [25] are low-latency AC protocols, i.e., they immediately forward messages. While Tor and Hornet do not produce asymptotically more than a constant amount of both bandwidth overhead and latency overhead and thus cannot provide strong anonymity, Herd produces dummy traffic linearly proportional to the number of users (bandwidth overhead $\beta \in \theta(N/N)$), thus the trilemma does *not* exclude strong anonymity for Herd.

Riposte [50] uses secure multiparty computation and a variant of PIR to implement an anonymous bulletin board. Riposte operates in epochs and for each epoch the set of users is public. Hence, Riposte is expected to be run with long epochs to maximize the number of users that participate in an epoch, which leads us to estimating the latency overhead to be $\ell \in \theta(N)$. To counter traffic analysis attacks, Riposte clients send constant dummy traffic, resulting in a bandwidth overhead of $\beta \in \theta(N/N)$. Thus, the trilemma does not exclude strong anonymity for Riposte.

Vuvuzela [20] is a mix-net that is tailored towards messengers. Clients communicate by depositing their encrypted messages in one of the mix net nodes. To achieve strong resistance against compromised servers, Vuvuzela takes a path through all servers, resulting in a latency overhead of $\ell \in \theta(K)$




Fig. 6. Asymptotic latency overhead (ℓ) and bandwidth overhead (β) together with the “area of impossibility” where $2\ell\beta \leq 1 - \frac{1}{\text{poly}(\eta)}$. We portray protocols as dots depending on their choices for ℓ and β . Technically, if we use Theorem 7, we β is replaced by $p = \beta + p'$, where p' is the rate at which users send messages. This graph assumes N is ca. $\text{poly}(\eta)$, the number of nodes K is ca. $\log \eta$. The threshold for Threshold Mix $T = 1$ and for Threshold Mix_{sec} $T = N = \text{poly}(\eta)$. In the graph, both the axes are approximately in logarithmic scale. (For a more accurate visual representation we refer the readers to Appendix C and [51].)

(for K servers). Additionally, Vuvuzela utilizes constant traffic, leading to a bandwidth overhead of $\beta \in \theta(N/N)$, and has the potential for strong anonymity.

Riffle [21] uses a verifiable mix-net. Just as Vuvuzela, Riffle also chooses paths that traverse all K servers, leading to $\ell \in \theta(K)$ and if we assume $K \in \theta(\log(\eta))$, we get $\ell \in \theta(\log(\eta))$. We assume that the clients send dummy traffic up to a constant rate (depending on the user’s sending rate p'), so we have $\beta \in \theta(N/N)$ and the potential for strong anonymity.

In a **threshold mix net**, each of the K mix servers waits until it received up to a threshold T many messages before relaying the messages to the next mix, resulting in $\ell \in \theta(T \times K)$. Threshold mixes [14] do not provide strong anonymity unless their threshold T is of the order of the number of users N . As such a large threshold are impractical for a large number of users, we judge it impossible to achieve strong anonymity for practical of Threshold mixes.

Loopix [24] is a mix net that combines exponentially distributed delays at each mix-node and dummy messages from each user. Ignoring so-called loop messages (meant to counter active attacks), Loopix naturally enforces our unsynchronised user distribution: the rate at which Loopix clients send messages is the sum of a dummy-message rate (β) and a payload message rate (p'), which are system parameters. We assume that the path lengths in Loopix’ stratified topology is \sqrt{K} with the number of nodes $K \in \theta(\log(\eta))$. If $\beta + p' \geq 1/\sqrt{\eta}$, and if every hop introduces an expected delay of $\ell' \geq \frac{\sqrt{\eta}}{\sqrt{K}}$, the expected latency overhead is $\ell = \sqrt{K} \times \ell'$, in particular $\ell \in \theta(\sqrt{\log(\eta)})$. We get $(p' + \beta)\ell = \frac{1}{\sqrt{\eta}} \times \sqrt{\eta} = 1$ and the trilemma does not exclude strong anonymity for Loopix.

In **AC protocols based on DC-nets** [15], [18] each party broadcasts either a dummy or real message in every round to every other party. As our bandwidth overhead only counts dummy-message rates, it does not capture the broadcast, thus $\beta \in \theta(N/N)$. DC-nets use a combination operation (e.g., an XOR) that causes dummy messages to cancel out. Then, all parties output the resulting bitstring. If only one real message is sent, the bitstring equals this message. As Theorem 7 assumes a synchronized user distribution, in each round only one party sends a message, thus our model treats ℓ as $\ell \in \theta(1)$.

The **Dissent-AT** [22] scheme (the AnyTrust-variant of Dissent) improves on the performance of DC-nets by relying on dedicated servers. Instead of broadcasting to every other client, clients in Dissent-AT send these messages to at least one of these dedicated servers. These servers then perform a DC-net communication round. Abstracting from an initial setup phase and only counting the client-messages, Dissent-AT has $\beta \in \theta(N/N)$ for the clients (assuming that each client communicates to one server), and $\ell \in \theta(1)$.

Dicemix [16] is a peer-to-peer AC protocol that is based on the DC-net approach. While Dicemix includes a self-healing mechanism that leads to $4 + 2f$ communication rounds for one message if f peers are malicious, this mechanism does not kick in if all peers are honest, leading to only 4 communication rounds, resulting in $\ell \in \theta(1)$. As every party sends a message in every round $\beta \in \theta(N/N)$.

XI. CONCLUSION AND FUTURE WORK

This paper proves the anonymity trilemma: strong anonymity, low bandwidth, low latency—choose two! We derive necessary constraints for sender anonymity and recipient anonymity, and thereby presents necessary constraints that are crucial for understanding bi-directional anonymous communication: sender anonymity for hiding the sender and recipient anonymity for hiding the recipient of a message.

For future work, we plan to extend the work in four natural directions: (i) derive tighter bounds by using more sophisticated attackers, (ii) derive bounds for other anonymity notions (e.g., unlinkability and relationship anonymity), (iii) extend the protocol mode with a notion of a throughput limitation, (iv) relax the requirement that messages are sent with certainty and allow for unreliable channels. For example, for the first direction, we plan to take the same steps as outlined in Section II-B, i.e., to formulate an invariant, to construct a protocol optimal w.r.t. this invariant, and then to compute the advantage of the more sophisticated attacker against this protocol.

ACKNOWLEDGMENTS

We thank the reviewers for their valuable comments. This work has been partially supported by the Zurich Information Security Center (ZISC), the European Commission through H2020-DS-2014-653497 PANORAMIX, the EPSRC Grant EP/M013-286/1, and the National Science Foundation (NSF) under grant CNS-1719196.

REFERENCES

- [1] “The Tor Project,” <https://www.torproject.org/>, accessed in Nov 2017.
- [2] A. Johnson, C. Wacek, R. Jansen, M. Sherr, and P. Syverson, “Users get routed: Traffic correlation on tor by realistic adversaries,” in *Proc. ACM SIGSAC conference on Computer & communications security 2013*, 2013, pp. 337–348.
- [3] L. Øverlier and P. F. Syverson, “Locating Hidden Servers,” in *Proc. 27th IEEE Symposium on Security and Privacy*, 2006, pp. 100–114.
- [4] S. J. Murdoch and G. Danezis, “Low-cost traffic analysis of Tor,” in *Proc. IEEE Symposium on Security and Privacy 2005*, 2005.
- [5] K. S. Bauer, D. McCoy, D. Grunwald, T. Kohno, and D. C. Sicker, “Low-resource routing attacks against tor,” in *Proc. 6th ACM Workshop on Privacy in the Electronic Society (WPES)*, 2007, pp. 11–20.
- [6] Y. Sun, A. Edmundson, L. Vanbever, O. Li, J. Rexford, M. Chiang, and P. Mittal, “RAPTOR: Routing attacks on privacy in Tor,” in *Proc. 24th USENIX Security Symposium*, 2015.
- [7] R. Jansen, F. Tschorsch, A. Johnson, and B. Scheuermann, “The sniper attack: Anonymously deanonymizing and disabling the Tor network,” in *Proc. Network and Distributed Security Symposium - NDSS '14*, 2014.
- [8] Y. Gilad and A. Herzberg, “Spying in the Dark: TCP and Tor Traffic Analysis,” in *Proc. 12th Privacy Enhancing Technologies Symposium (PETS 2012)*, 2012.
- [9] The Tor Blog, “One cell is enough to break Tor’s anonymity,” <https://blog.torproject.org/blog/one-cell-enough>, accessed Nov 2017.
- [10] R. Dingledine, N. Mathewson, and P. Syverson, “Tor: The Second-Generation Onion Router,” in *Proc. 13th USENIX Security Symposium (USENIX)*, 2004, pp. 303–320.
- [11] S. Chakravarty, M. V. Barbera, G. Portokalidis, M. Polychronakis, and A. D. Keromytis, “On the effectiveness of traffic analysis against anonymity networks using flow records,” in *Proc. 15th International Conference on Passive and Active Measurement*, 2014, pp. 247–257.
- [12] N. Gelernter and A. Herzberg, “On the limits of provable anonymity,” in *Proc. Workshop on Privacy in the Electronic Society (WPES 2013)*, 2013, pp. 225–236.
- [13] A. Hevia and D. Micciancio, “An indistinguishability-based characterization of anonymous channels,” in *Proc. Eighth International Symposium on Privacy Enhancing Technologies (PETS 2008)*, N. Borisov and I. Goldberg, Eds., 2008, pp. 24–43.
- [14] A. Serjantov, R. Dingledine, and P. Syverson, “From a trickle to a flood: Active attacks on several mix types,” in *5th Information Hiding Workshop (IH 2002)*, 2003, pp. 36–52.
- [15] D. Chaum, “The dining cryptographers problem: Unconditional sender and recipient untraceability,” *Journal of Cryptology*, vol. 1, no. 1, pp. 65–75, 1988.
- [16] T. Ruffing, P. Moreno-Sanchez, and A. Kate, “P2P Mixing and Unlinkable Bitcoin Transactions,” in *Proc. 25th Annual Network & Distributed System Security Symposium (NDSS)*, 2017.
- [17] H. Corrigan-Gibbs and B. Ford, “Dissent: Accountable Anonymous Group Messaging,” in *Proc. 17th ACM Conference on Computer and Communication Security (CCS)*, 2010, pp. 340–350.
- [18] P. Golle and A. Juels, “Dining cryptographers revisited,” in *Proc. of Eurocrypt 2004*, 2004.
- [19] H. Corrigan-Gibbs, D. I. Wolinsky, and B. Ford, “Proactively Accountable Anonymous Messaging in Verdict,” in *Proc. 22nd USENIX Security Symposium*, 2013, pp. 147–162.
- [20] J. van den Hooff, D. Lazar, M. Zaharia, and N. Zeldovich, “Vuvuzela: Scalable private messaging resistant to traffic analysis,” in *Proc. 25th ACM Symposium on Operating Systems Principles (SOSP 2015)*, 2015.
- [21] A. Kwon, D. Lazar, S. Devadas, and B. Ford, “Riffle: An Efficient Communication System With Strong Anonymity,” in *Proc. Privacy Enhancing Technologies Symposium (PETS 2016)*, 2016, pp. 115–134.
- [22] D. I. Wolinsky, H. Corrigan-Gibbs, B. Ford, and A. Johnson, “Dissent in Numbers: Making Strong Anonymity Scale,” in *10th USENIX Symposium on Operating Systems Design and Implementation (OSDI 12)*, 2012, pp. 179–182.
- [23] S. Le Blond, D. Choffnes, W. Zhou, P. Druschel, H. Ballani, and P. Francis, “Towards Efficient Traffic-analysis Resistant Anonymity Networks,” in *Proc. ACM SIGCOMM 2013*, 2013, pp. 303–314.
- [24] A. Piotrowska, J. Hayes, T. Elahi, S. Meiser, and G. Danezis, “The loopix anonymity system,” in *Proc. 26th USENIX Security Symposium*, 2017.
- [25] S. Le Blond, D. Choffnes, W. Caldwell, P. Druschel, and N. Merritt, “Herd: A Scalable, Traffic Analysis Resistant Anonymity Network for VoIP Systems,” in *Proc. ACM Conference on Special Interest Group on Data Communication (SIGCOMM 2015)*, 2015, pp. 639–652.
- [26] M. Backes, A. Kate, P. Manoharan, S. Meiser, and E. Mohammadi, “AnoA: A Framework For Analyzing Anonymous Communication Protocols,” in *Proc. 26th IEEE Computer Security Foundations Symposium (CSF 2013)*, 2013, pp. 163–178.
- [27] M. Backes, A. Kate, P. Manoharan, S. Meiser, and E. Mohammadi, “AnoA: A Framework For Analyzing Anonymous Communication Protocols,” *Journal of Privacy and Confidentiality (JPC)*, vol. 7(2), no. 5, 2016.
- [28] K. Jensen, *Colored Petri Nets. Basic Concepts, Analysis Methods and Practical Use.*, 1997, vol. 3.
- [29] W. Reisig, *Primer in Petri Net Design*, 1st ed., 1992.
- [30] L. M. Kristensen, S. Christensen, and K. Jensen, “The practitioners guide to coloured petri nets,” *International Journal on Software Tools for Technology Transfer (STTT)*, vol. 2, no. 2, pp. 98–132, 1998.
- [31] T. K. Srikanth and S. Toueg, “Simulating authenticated broadcasts to derive simple fault-tolerant algorithms,” *Distributed Computing*, vol. 2, no. 2, pp. 80–94, 1987.
- [32] R. Gennaro, M. O. Rabin, and T. Rabin, “Simplified VSS and fact-track multiparty computations with applications to threshold cryptography,” in *Proc. ACM PODC*, 1998, pp. 101–111.
- [33] D. Das, S. Meiser, E. Mohammadi, and A. Kate, “Anonymity trilemma: Strong anonymity, low bandwidth, low latency—choose two,” *Cryptology ePrint Archive*, Report 2017/954, 2017, <https://eprint.iacr.org/2017/954>.
- [34] J. Feigenbaum, A. Johnson, and P. Syverson, “A probabilistic analysis of onion routing in a black-box model,” in *Proc. Workshop on Privacy in the Electronic Society (WPES 2007)*, 2007.
- [35] D. Wikström, “A Universally Composable Mix-Net,” in *Proc. 1st Theory of Cryptography Conference (TCC)*, 2004, pp. 317–335.
- [36] J. Camenisch and A. Lysyanskaya, “A formal treatment of onion routing,” in *Proc. CRYPTO 2005*, 2005, pp. 169–187.
- [37] N. Kiyavash, A. Houmansadr, and N. Borisov, “Multi-flow Attacks Against Network Flow Watermarking Schemes,” in *Proc. 17th USENIX Security Symposium*, 2008.
- [38] S. Oya, C. Troncoso, and F. Pérez-González, “Do dummies pay off? limits of dummy traffic protection in anonymous communications,” in *Proc. 14th Privacy Enhancing Technologies Symposium (PETS 2014)*, 2014.
- [39] G. Danezis, “Statistical disclosure attacks: Traffic confirmation in open environments,” in *Proc. Security and Privacy in the Age of Uncertainty (SEC2003)*, 2003, pp. 421–426.
- [40] G. Danezis and A. Serjantov, “Statistical disclosure or intersection attacks on anonymity systems,” in *Proc. 6th Information Hiding Workshop (IH 2004)*, 2004.
- [41] M. J. Freedman, K. Nissim, and B. Pinkas, “Efficient private matching and set intersection,” in *Proc. EUROCRYPT 2004*, 2004.
- [42] F. Pérez-González and C. Troncoso, “Understanding statistical disclosure: A least squares approach,” in *Proc. 12th Privacy Enhancing Technologies Symposium (PETS 2012)*, 2012, pp. 38–57.
- [43] M. Backes, P. Manoharan, and E. Mohammadi, “TUC: Time-sensitive and Modular Analysis of Anonymous Communication,” *IACR ePrint Archive Report 2013/664*, 2013, <http://www.infsec.cs.uni-saarland.de/mohammadi/paper/tuc.pdf>.
- [44] R. Dingledine, N. Mathewson, and P. Syverson, “Tor: The second-generation onion router,” in *Proc. 13th USENIX Security Symposium*, 2004.
- [45] D. Dolev, R. Reischuk, and H. R. Strong, “Early stopping in byzantine agreement,” *J. ACM*, vol. 37, no. 4, pp. 720–741, 1990.
- [46] D. Kesdogan, J. Egner, and R. Büschkes, “Stop-and-go MIXes: Providing probabilistic anonymity in an open system,” in *Proc. Information Hiding Workshop (IH 1998)*, 1998.
- [47] G. Danezis, C. Diaz, C. Troncoso, and B. Laurie, “Drac: An architecture for anonymous low-volume communications,” in *Proc. 10th Privacy Enhancing Technologies Symposium (PETS 2010)*, 2010.
- [48] P. Mittal, M. Wright, and N. Borisov, “Pisces: Anonymous communication using social networks,” in *Proc. 20th Network and Distributed System Security Symposium (NDSS 2013)*, 2013.
- [49] C. Chen, D. E. Asoni, D. Barrera, G. Danezis, and A. Perrig, “HOR-NET: High-speed onion routing at the network layer,” in *Proc. ACM*

Conference on Computer and Communications Security (CCS), 2015, pp. 1441–1454.

- [50] H. Corrigan-Gibbs, D. Boneh, and D. Mazières, “Riposte: An anonymous messaging system handling millions of users,” in *Proc. 36th IEEE Symposium on Security and Privacy (S&P 2015)*, 2015, pp. 321–338.
- [51] Anonymity Trilemma Project Webpage, “Anonymity trilemma: Strong anonymity, low bandwidth overhead, low latency overhead—choose two,” <https://freedom.cs.purdue.edu/projects/anonymity/trilemma/>.

TABLE I

Latency vs. bandwidth vs. strong anonymity of AC protocols, with the number of protocol-nodes K , number of clients N , and message-threshold T , expected latency ℓ' per node, dummy-message rate β .

Protocol	Latency	Bandwidth	Strong Anonymity
Tor [10]	$\theta(1)$	$\theta(1/N)$	impossible
Hornet [49]	$\theta(1)$	$\theta(1/N)$	impossible
Herd [25]	$\theta(1)$	$\theta(N/N)$	possible
Riposte [50]	$\theta(N)$	$\theta(N/N)$	possible
Vuvuzula [20]	$\theta(K)$	$\theta(N/N)$	possible
Riffle [21]	$\theta(K)$	$\theta(N/N)$	possible
Threshold mix [14]	$\theta(TK)$	$\theta(1/N)$	impossible*
Loopix [24]	$\theta(\sqrt{K}\ell')$	$\theta(\beta)$	possible
DC-Net [15], [18]	$\theta(1)$	$\theta(N/N)$	possible
Dissent-AT [22]	$\theta(1)$	$\theta(N/N)$	possible
DiceMix [16]	$\theta(1)$	$\theta(N/N)$	possible

* if T in $o(\text{poly}(\eta))$

APPENDIX A

PROTOCOL MODEL REVISITED

A. Validity of the Protocol Model (Contd.)

Lemma 2. *Let Π be a protocol $\in M$ with K parties with parameters β and ℓ . Then: 1) Messages are delivered within ℓ steps. 2) The protocol adds (for the unsynchronised case on average) a maximum of β noise messages per user per round. 3) Whenever a party in $S \cup P$ sends a message to another party in $P \cup R$, the adversary learns that and in which round this happens. 4) For every message that leaves the network (received by R), the adversary additionally learns whether the message is the target message. 5) For every compromised party, the adversary learns the mapping between the input messages and the output messages.*

Proof. Let Π be a protocol $\in M$ with K parties with parameters β and ℓ . Part (2) of the Lemma holds, since we restrict the user distributions accordingly and since the none of the transitions in the petri-net can create more tokens within the network than it consumes from its input place.

We show the part (1) of the lemma via structural induction over fired transitions of the petri net. We additionally add to the induction invariant that all tokens that are not in S have a timestamp for their next transition of $ts = 1$ and a remaining time of $t_r > 0$ and there are at least t_r rounds left in which the token can be delivered.

Induction base: The protocol is initialized and no transitions have happened. Thus, no messages have been sent so far, i.e., there is no message that has not been delivered within ℓ steps. The only transition that can fire is T_S and for $\ell > 0$, the message introduced into the network in this way does not need to be delivered already ($0 < t_r = \ell$). Moreover, T_S sets the timestamp of this message token to $ts = 1$

Induction step: Let tr be any execution trace s.t. the induction invariant is satisfied and let t be an arbitrary possible transition that extends tr to $tr :: t$.

We distinguish two cases for t : In case t is T_S , it consumes a token from P_S and puts this token into a place P_i and, by definition we have $t_r > 0$ and $ts = 1$. Otherwise, the transition is T_{P_i} for some i and consumes a token from P_i accordingly. By the induction invariant, the token has $t_r > 0$. If this token has $t_r - 1 = 0$, the transition delivers the token to R . Otherwise, t decreases t_r by one (thus fulfilling the condition that there are at least t_r rounds left in which the token can be delivered) and sets $ts = 1$. Since every token in any place P_i needs to be consumed in every round, the protocol delivers every message in at most ℓ steps.

Other parts of the lemma: By definition of our petri net, whenever a transition fires, an element (t, r) is placed into Tokens, containing the public fields of t , such as $t.\text{prev}$ and $t.\text{next}$, as well as the current round number r , which fulfills part (3). Moreover, whenever the transition places the token in R , the adversary can additionally see the field $t.\text{msg}$ and no transition can change the field msg , which allows the adversary to effectively tag and recognize the challenge message and thus fulfills part (4). Finally, if any party P_i is compromised, P_i does not modify the unique (and otherwise freshly sampled) field $t.\text{ID}_t$, which allows the adversary to map incoming and outgoing messages.

Since the transitions discussed here are the only way for messages to be sent to a recipient, the model correctly enforces the conditions from the lemma. \square

B. Expressing Protocols in the petri net model

Modeling DC net. Here we show how to model an actual DC net type protocol using our petri net model M as defined in Section IV. Specifically we pick up the *short DC net* protocol proposed by Golle and Juels [18], and present M_{DC} which models the aforementioned protocol.

We model a DC net protocol with N participants, where $S = P$, $|S| = |P| = N$. We denote the parties with P_1, \dots, P_N . The protocol can be denoted by $\Pi_{DC} = \{\text{paramgen}, \text{keydist}, \text{post}, \text{verify}, \text{extract}\}^7$ - as described below.

- *paramgen*: In prot_{DC} , *paramgen* is executed by a trusted entity and the output is published. Since we are mainly interested in the anonymity game, we consider that *paramgen* step is executed by our honest challenger and happens outside the protocol run, and the output is globally known (to all the transitions T_{P_i}).

- *keydist*: using the output of *paramgen*, this step yields for each party P_i a private key x_i and a corresponding public key y_i . In prot_{DC} , the above key generation part is done by a trusted entity, and hence we consider that it is done by our honest challenger and for each party P_i the public-private keypair x_i, y_i is already known to the corresponding transition

⁷Since we are mainly interested in the anonymity property, we don't need to model the part of the protocol where the protocol parties reconstructs the keys in case of a failure. But it is easy to extend M_{DC} to include that step by adding one more round to the current model.

function T_{P_i} . As part of protocol each party P_i publishes its public key y_i . Additionally, each party P_j receives from P_i a share of private key $x_{i,j}$ and a share of public key $y_{i,j}$, where the keys are shared in a (k, N) threshold manner for a parameter $k \leq N$.

- *post*: Each player P_i generates a vector of random pads $W_i = \{W_i(1), W_i(2), \dots, W_i(N)\}$ ⁸ using x_i . Π_{DC} does not handle *collisions*, instead assumes that the players decide their positions by a consensus protocol. Similarly our model assumes that each party P_i knows its position, and assume the position is q_i (but not known to the adversary). Then each player P_i computes the vector V_i such that $V_i(w) = W_i(w)$ for all $w \neq i$ and $V_i(w) = W_i(w) \oplus m_i$ for $w = q_i$, where m_i is the message of P_i . Also, each player P_i computes $\sigma_i = \{\sigma_i(1), \sigma_i(2), \dots, \sigma_i(N)\}$, where σ_i includes the identity of player P_i and a proof of valid formatting of V_i . Then P_i publishes both the vectors V_i and σ_i . Our model assumes the pair $(V_i(w), \sigma_i(w))$ for each position w as a single message, where $V_i(w)$ is a message content and $\sigma_i(w)$ becomes a part of *meta* field. For each position w player P_i generates one such message, and publishes the message to all other players.

- *verify* and *extract* are local computations after a party P_i receives messages from all other parties.

Although the protocol model assumes that the adversary can not read the contents of any message, here we shall model Π_{DC} along with its cryptographic primitives to demonstrate the expressiveness of our model. Alternatively, to get rid of all the cryptographic primitives, the parties can send a dummy message ($= 0$) whenever $V_i(w) = W_i(w)$, and the actual message m_i whenever $V_i(w) \neq W_i(w)$.

As per our anonymity definition in Section III, we assume that up to $(N-2)$ users can be compromised, which necessarily makes up to $(N-2)$ protocol parties compromised. The adversary chooses two challenge users, and one of them sends the challenge message depending on the challenge bit b . All other $(N-1)$ users send dummy messages.

In M_{DC} we model Π_{DC} as a two round protocol. The challenger sets the initial configuration of the petri-net with the messages to be sent by each party. In the first round, each party P_i sends two kinds messages: (1) publishes the public key message y_i and (2) sends share of the public-private keypair $(x_{i,j}, y_{i,j})$ to P_j for all $j \neq i$. Here, one party can publish a message to $(N-1)$ other parties by sending $(N-1)$ separate messages. In the second round, each party P_i publishes N messages: one message for each position, only one of them contains his own message. After second round, every party receives messages from every other party, and then does local computations to verify and extract the original messages.

For Π_{DC} , we do not actually need a separate recipient R in Π_{DC} , if we make $\mathcal{R} = \mathcal{P}$. But, to be consistent with M , in M_{DC} we keep a separate recipient. In the second round whenever a party P_i publishes a message, P_i also sends a copy to R . This easily models the fact that the adversary knows

⁸The anonymity game does not include multiple sessions. Also, in our model all the N players participate in a protocol run.

whenever a message is published, but avoids the complication of modeling a subset of compromised recipients.

The *meta* fields of the tokens contains the following sub-fields: (1) stage, (2) position, (3) sigma. *stage* can have three possible values identifying three possible cases: (1) public key distribution, (2) share of the public-private keypair, (3) message. Using *stage* subfield, any party in the protocol recognizes if the message is part of *keydist* messages, or part of *post* messages. When the value of *stage* is *message*, the user posts $V_i(w)$, and *position* takes the value of w . *sigma* includes the identity of the sender and a proof of computation whenever necessary. *sigma* fields helps in the *verify* stage, we avoid the details here.

If we want to analyze the user distribution for Π_{DC} , we do not count the first round since it is used only for key exchange. Note that, if we get rid of the cryptographic primitives, we do not require the first round.

Modeling Tor. Since Tor does not operate in rounds, embedding it into our model is not straight forward. Suppose, a Tor node takes at least x milliseconds to process a message when it receives a message, and it takes at least y milliseconds for a message to travel from one node to the next node over a network link. Then we define *one round* as $x+y$ milliseconds. We assume a perfect condition where each node takes exactly equal computation time for one message, and each link has exactly same delay. In the real world, delays and computation times are less stable, but can be estimated by an adversary. Instead of analyzing this, we instead allow the messages to remain within the node for the respective time.

Tor nodes and recipients are separate entities and hence, \mathcal{S} , \mathcal{P} and \mathcal{R} are mutually exclusive. Whenever a Tor node receives a message, the node immediately processes and forwards that message to the next node or recipient. We can either model the latency overhead ℓ of Tor by estimating the time messages spend within the network that exceeds the (minimal) round length $x+y$ from above, or we set it to the number of hops, i.e., $\ell = 3$. In either case, we assume that ℓ does not increase with η and thus get a latency overhead $\ell \in O(1)$. For analyzing Tor with a variable number h of hops, we can instead set $\ell = h$. When a party P_i receives a message, T_{P_i} can retrieve the next hop from the meta field of the message. Since Tor does not add any noise messages, the bandwidth overhead is $\beta = 0$.

APPENDIX B DELAYED PROOFS

Proof of Claim 1. If the set T is empty, then S_{1-b} is empty as well. However, by construction of our protocol mode, the set S_b is always non-empty. Consequently, the adversary \mathcal{A}_{paths} will output b and thus win with probability 1. If T is not empty, the following cases can occur:

- 1) The challenge message never passes through an honest node: In this case, the field ID_t of the message never changes for the tokens. Thus, S_b will have exactly one element, and S_{1-b} will be an empty set, and consequently \mathcal{A}_{paths} wins.

- 2) The challenge message passes through one or more honest nodes at times t' , such that $t' < \min(T)$, but not

afterwards. Following the same reasoning as above, we see that paths before $\min(T)$ can be ambiguous, but none of them leads to u_{1-b} . Hence, S_b can have multiple elements, but S_{1-b} will still be an empty set. Thus, \mathcal{A}_{paths} wins.

3) The challenge message passes through an honest node at time t' with $t' \geq \min(T)$. In this case, the invariant is true.

In all of the above mentioned cases either the invariant is true, or the adversary wins with probability 1. \square

Proof of Claim 2. Let u_0, u_1 be the challenge users and let b be the challenge bit and let r be the round in which the challenge message is delivered to the recipient. We discuss both parts of the invariant separately:

(i) The set T is not empty. Since by definition, T is the set of messages sent by u_{1-b} , messages originating in any party not in \mathcal{S} do not influence T . Moreover, any message sent by u_{1-b} in a round previous to $r - \ell$ does not influence T either. Thus, noise messages staying in the protocol for more than ℓ rounds, do not improve the probability of T being not empty.

(ii) The challenge message passes through at least one honest node at some time t' such that, $t' \in \{\min(T), \dots, r - 1\}$. Obviously this second part of the invariant does not depend on any noise message. \square

Proof of Claim 3. We want to prove our claim by contradiction. Suppose, Π_{ideal} is not the best protocol. That means, there exists a protocol Π_{new} , which satisfies Invariant 1 with a higher probability than Π_{ideal} , against the adversary \mathcal{A}_{paths} .

Now we construct a new protocol Π_{hybrid} , which exactly follows the strategy of Π_{ideal} with one exception: for a given message Π_{hybrid} selects the time delay t same as Π_{new} , instead of querying it from oracle \mathcal{O} . Suppose, the challenge message is delivered to the recipient at round r . Given the set $\{\min(T), \dots, r - 1\}$, the ideal strategy for ensuring that at least one honest party is on the path of the challenge message is to ensure that as many distinct parties as possible are on this path. Also, given the time delay t , the value of $\min(T)$ is independent of the protocol, since protocols in \mathcal{M} are oblivious to the challenge users and the challenge message. Hence, Π_{hybrid} has a probability of satisfying Invariant 1 at least as high as Π_{new} .

Now, if we compare Π_{hybrid} and Π_{ideal} : they follow the same strategy. But Π_{ideal} picks the time delay t for any message from oracle \mathcal{O} (except for messages from u_{1-b}) such that t is *optimal*. The time delay t can be picked for each message independent of the time delays of other messages. Hence, the value of t received from oracle \mathcal{O} for the challenge message is optimal. Hence, Π_{ideal} satisfies Invariant 1 with probability at least as high as Π_{hybrid} . Thus, Π_{new} does not satisfy Invariant 1 with a higher probability than Π_{ideal} . \square

Proof of Claim 4. If the Invariant is true, the challenge message passes through an honest party at t' , such that $t' > \min(T)$. Hence, there is at least one message (noise or original message) from u_{1-i} which visits the same honest party together with the challenge message (Π_{ideal} ensures that all messages are always kept together until they are delivered).

That ensures that in addition to $S_b \neq \emptyset$, we also have $S_{1-b} \neq \emptyset$ and thus \mathcal{A}_{paths} outputs a random bit (and has an advantage of zero). \square

Proof of Theorem 2. For strong anonymity, we require: $\delta(\eta) = \text{neg}(\eta)$, and we know that for Π_{ideal} we have: $\delta(\eta) \geq 1 - f_\beta(\ell) = \left(\frac{N - \ell - \beta N \ell}{N - 1}\right) \geq \left(\frac{N - \ell - \beta N \ell}{N}\right) \geq 1 - \frac{\ell}{N} - \beta \ell$. We assume for contradiction that there is a protocol limited by ℓ and β such that $2\ell\beta < 1 - \epsilon(\eta)$ that still achieves strong anonymity. Since $\delta(\eta) = \text{neg}(\eta)$, we know that $\epsilon(\eta) > \delta(\eta)$.

$$\begin{aligned} \epsilon(\eta) > \delta(\eta) &\implies \epsilon(\eta) > 1 - \frac{\ell}{N} - \beta \ell \\ &\implies \epsilon(\eta) > 1 - \frac{\ell}{N} - \frac{1}{2}(1 - \epsilon(\eta)) \\ &\iff 2\ell > N(1 - \epsilon(\eta)) \stackrel{N\beta \geq 1}{\implies} 2\ell\beta > 1 - \epsilon(\eta) \end{aligned}$$

The above contradicts the assumption that $2\ell\beta < 1 - \epsilon(\eta)$.

Note: In case $\beta N < 1$, no noise messages are allowed per round (i.e., $\beta = 0$) and thus $\delta(\eta) \geq 1 - \ell/N$, which is not negligible unless $\ell = N$, since $N = \text{poly}(\eta)$. \square

Proof of Theorem 4. When $c > \ell$: $\delta \geq 1 - \left[1 - \frac{\binom{c}{\ell}}{\binom{K}{\ell}}\right] f_\beta(\ell)$.

For δ to become *neg*(η), we need both $[1 - \frac{\binom{c}{\ell}}{\binom{K}{\ell}}]$ and $f_\beta(\ell)$ to become overwhelming. From Theorem 2 and Theorem 1, we know that $2\ell\beta > 1 - \text{neg}(\eta)$ is a necessary condition for $f_\beta(\ell)$ to become overwhelming. Now, we are left with the factor $[1 - \frac{\binom{c}{\ell}}{\binom{K}{\ell}}]$. This can become overwhelming iff $[\frac{\binom{c}{\ell}}{\binom{K}{\ell}}]$ becomes negligible. We know that $K > c \geq \ell$ and $K \in \text{poly}(\eta)$. Hence, for some constant x ,

$$\begin{aligned} \frac{c - \ell}{K - \ell} > \frac{1}{\eta^x} &\iff \left(\frac{c - \ell}{K - \ell}\right)^\ell > \left(\frac{1}{\eta^x}\right)^\ell \\ &\implies \frac{c(c-1)\dots(c-\ell)}{K(K-1)\dots(K-\ell)} > \left(\frac{c - \ell}{K - \ell}\right)^\ell > \left(\frac{1}{\eta^x}\right)^\ell \\ &\iff \frac{\binom{c}{\ell}}{\binom{K}{\ell}} > \left(\frac{1}{\eta^x}\right)^\ell. \end{aligned}$$

For any $\ell \in \mathcal{O}(1)$, $(1/\eta^x)^\ell$ is non-negligible. \square

Proof of Theorem 5. When $c < \ell$:

$$\delta \geq 1 - \left[1 - 1/\binom{K}{c}\right] f_\beta(c) - f_\beta(\ell - c).$$

First consider the factor $[1 - 1/\binom{K}{c}]$. Since $K = \text{poly}(\eta)$ and $c = \text{constant}$, $[1/\binom{K}{c}]$ can never be negligible. And thus, $[1 - 1/\binom{K}{c}]$ can never be overwhelming. So, $[1 - 1/\binom{K}{c}]f_\beta(c)$ can never be overwhelming as well, since $f_\beta(c) \leq 1$.

Now, let's consider $f_\beta(\ell - c)$ and $f_\beta(c)$. Note that, these two factors represent the probabilities of two dependent but mutually exclusive events, and hence $f_\beta(c) + f_\beta(\ell - c) \leq 1$. And we already know that $[1 - 1/\binom{K}{c}]$ can never be overwhelming. Thus, the only way δ can become negligible is if $f_\beta(\ell - c)$ becomes overwhelming. Note that, if $a + b \leq 1$ and $c < 1$, the only way $ac + b = 1$ is possible if $b = 1$.

Now we can follow exactly the same procedure as in the proof of Theorem 2 to say: $f_\beta(\ell - c)$ can not become overwhelming if $2(\ell - c)\beta < 1 - \epsilon(\eta)$. \square

Proof of Theorem 7. We know $0 \leq E \leq 1/2$. When $2\mu \leq N$,

$$\begin{aligned} \delta &\geq (1-E)(1-2f_p(\ell)) \geq 1/2 \left(2(1-p)^\ell - 1 \right) \\ &\geq 1/2(2(1-\ell p) - 1) = 1/2(1-2\ell p). \end{aligned}$$

Thus, if $2\ell p < 1 - \epsilon(\eta)$,

$$\begin{aligned} 2\ell p < 1 - \epsilon(\eta) &\iff 1 - 2\ell p > \epsilon(\eta) \\ &\implies \delta > 1/2 \times \epsilon(\eta) = \text{non-negligible}. \end{aligned}$$

Thus, when $2\mu \leq N$, a necessary condition for δ to become negligible is $2\ell p > 1 - \text{neg}(\eta)$.

When $2\mu > N$, using $\mu = N(1 - (1-p)^\ell)$ we get:

$$\begin{aligned} 2N(1 - (1-p)^\ell) > N &\implies (1-p)^\ell < 1/2 \\ \implies 1 - p\ell < 1/2 &\iff 2p\ell > 1. \quad \square \end{aligned}$$

Proof of Theorem 8. Let $X^{(i)}(x)$ and $X(x)$ be defined as in the proof for Theorem 6, where we replace the fixed length ℓ of the slice by a variable x . Using the Chernoff Bound on the random variable $X(x)$ calculate $\Pr[X(x) - \mu(x) \geq Na] \leq \exp(-2a^2N)$, and for $a = \frac{\mu(x)}{N}$, we define $E(x)$ as :

$$\begin{aligned} E(x) &= \Pr[X(x) \geq 2\mu(x)] \leq \exp(-2\mu(x)^2/N^2 \times N) \\ &\leq \exp(-2(1 - (1-p)^x)^2 N). \end{aligned}$$

Note that, similar to $X^{(i)}(x)$ and $X(x)$, $\mu(x)$ is also defined as in the proof for Theorem 6, but for a slice of variable length x . We denote the event that sender u_{1-b} sends at least one message in an interval of size x by $Y(x)$ and since all users are acting independently from each other we get for $j \in \{0, \dots, N\}$, $\Pr[Y(x)|X(x) = j] = 1 - \Pr[-Y|X(x) = j] = \frac{j}{N}$. Moreover, for any value of x with $2\mu(x) \leq N$,

$$\begin{aligned} \Pr[Y(x)] &= \Pr[X(x) \geq 2\mu(x)] \times \Pr[Y(x)|X(x) \geq 2\mu(x)] \\ &\quad + \Pr[X(x) < 2\mu(x)] \times \Pr[Y(x)|X(x) < 2\mu(x)] \\ &\leq \Pr[X(x) \geq 2\mu(x)] \times \Pr[Y(x)|X(x) = N] \\ &\quad + \Pr[X(x) < 2\mu(x)] \times \Pr[Y(x)|X(x) = 2\mu(x)] \\ &= E(x)\Pr[Y|X(x) = N] \\ &\quad + (1 - E(x))\Pr[Y|X(x) = 2\mu(x)] \\ &= E(x) \binom{N}{N} + (1 - E(x)) \binom{2\mu(x)}{N} \\ &= 1 - (1 - E(x)) (1 - 2(1 - (1-p)^x)). \end{aligned}$$

If $2\mu(x) > N$, we get with $f(x) = \min(\frac{1}{2}, 1 - (1-p)^x)$:

$$\begin{aligned} \Pr[Y(x)] &\leq E(x) + (1 - E(x)) \times 1 \leq 1 \\ &\leq 1 - (1 - E(x)) (1 - 2f(x)). \end{aligned}$$

Now, we calculate the probability of Invariant 1 being true, under our protocol Π_{ideal} and as in the proof for Theorem 3. We distinguish two cases depending on c and ℓ :

Case 1): $c > \ell$

$$\begin{aligned} &\Pr[\text{Invariant 1 is true}] \\ &\leq \Pr[-\text{Cmpr}(\ell)] \times \Pr[u_{1-b}.\text{sent}(r - \ell, r - 1)] \\ &= \Pr[-\text{Cmpr}(\ell)] \times \Pr[Y(\ell)] \\ &\leq \left[1 - \binom{c}{\ell} / \binom{K}{\ell} \right] \left[1 - (1 - E(\ell)) (1 - 2f_p(\ell)) \right]. \end{aligned}$$

By applying Markov's inequality on the random variable $X(x)$, we get $E(x) = \Pr[X(x) \geq 2\mu(x)] \leq \frac{1}{2}$. Thus, we derive for δ : $\delta \geq 1 - \left[1 - \binom{c}{\ell} / \binom{K}{\ell} \right] \left[\frac{1}{2} + f_p(\ell) \right]$.

Case 2): $c < \ell$. As for the proof of Theorem 3 we split this case into two sub-cases, depending on t and c .

Case 2a): $c < t$

$$\begin{aligned} &\Pr[\text{Invariant 1 is true}] \\ &\leq \Pr[u_{1-b}.\text{sent}(r - \ell, r - c)] + \Pr[-u_{1-b}.\text{sent}(r - \ell, r - c)] \\ &\quad \times \Pr[u_{1-b}.\text{sent}(r - c, r)] \times \Pr[-\text{Cmpr}(c)] \\ &= \Pr[Y(\ell - c)] + [1 - \Pr[Y(\ell - c)]] \Pr[Y(c)] \Pr[-\text{Cmpr}(c)] \\ &\leq [1 - (1 - E(\ell - c)) (1 - 2f_p(\ell - c))] \\ &\quad + [(1 - E(\ell - c)) (1 - 2f_p(\ell - c))] \\ &\quad \times [1 - (1 - E(c)) (1 - 2f_p(c))] \left[1 - 1 / \binom{K}{c} \right]. \end{aligned}$$

Thus, for the adversarial advantage δ we derive,

$$\begin{aligned} \delta &\geq 1 - \Pr[\text{Invariant 1 is true}] \\ &\geq 1 - [1 - (1 - E(\ell - c)) (1 - 2f_p(\ell - c))] \\ &\quad - [(1 - E(\ell - c)) (1 - 2f_p(\ell - c))] \\ &\quad \times [1 - (1 - E(c)) (1 - 2f_p(c))] \left[1 - \binom{c}{c} / \binom{K}{c} \right] \\ &= [(1 - E(\ell - c)) (1 - 2f_p(\ell - c))] \\ &\quad \times \left(1 - [1 - (1 - E(c)) (1 - 2f_p(c))] \left[1 - 1 / \binom{K}{c} \right] \right) \\ &\geq (1 - \left[\frac{1}{2} + f_p(\ell - c) \right]) \left(1 - \left[\frac{1}{2} + f_p(c) \right] \left[1 - 1 / \binom{K}{c} \right] \right). \end{aligned}$$

We again use Markov's inequality to replace $E(x)$ by $1/2$.

Case 2b): $t \leq c$

$$\begin{aligned} &\Pr[\text{Invariant 1 is true}] \\ &\leq \Pr[u_{1-b}.\text{sent}(r - \ell, r - c)] \times \Pr[-\text{Cmpr}(t)] \\ &\quad + \Pr[-u_{1-b}.\text{sent}(r - \ell, r - c)] \\ &\quad \times \Pr[u_{1-b}.\text{sent}(r - c, r)] \times \Pr[-\text{Cmpr}(c)] \\ &\leq \Pr[u_{1-b}.\text{sent}(r - \ell, r - c)] + \Pr[-u_{1-b}.\text{sent}(r - \ell, r - c)] \\ &\quad \times \Pr[u_{1-b}.\text{sent}(r - c, r)] \Pr[-\text{Cmpr}(c)] \end{aligned}$$

The above event expression is exactly same as the expression we had in the previous case ($t > c$). Thus, the rest of the calculations and bounds are exactly same as the previous case. \square

APPENDIX C

VISUAL 3D REPRESENTATIONS OF THE RESULTS

In the paper, we focus on lower-bound results for strong anonymity (or negligible δ values). However, our key Theorems 1, 3, 6 and 8 also offer lower bounds for non-negligible δ values, which can be of interest to several AC protocols.

On our project webpage [51], we visualize these lower bounds using interactive 3D surface plots. In particular, we plot the adversarial advantage $\delta \in [0, 1]$ as a function of β and ℓ . We encourage the readers to interact with these plots to better understand our results for non-negligible δ values.

Here, in Figures 7 to 10, we present and analyze four snapshots of those lower bound plots for the number of users $N = 10000$. The x -axis represents latency ℓ (ranging from 0 to 200), and the y -axis bandwidth overhead β (ranging from 0.0 to 0.04). But in Figure 9 and Figure 10, the y -axis actually represents total bandwidth $p = p' + \beta$ as in Theorem 7.

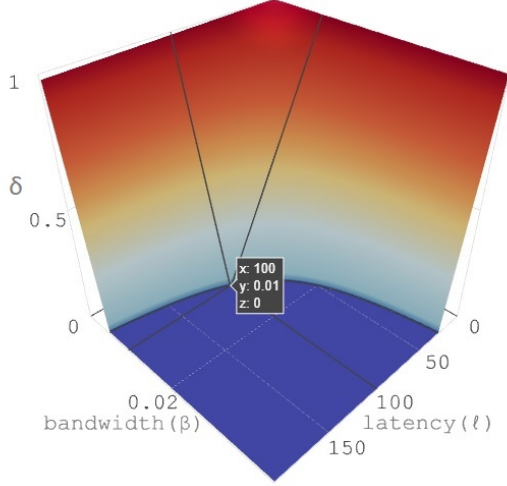


Fig. 7. Synchronized User Distribution with Non-compromising Adversaries. $z = 1 - f_\beta(\ell)$, where $f_\beta(x) = \min(1, ((x + \beta N x)/(N - 1)))$.




Fig. 9. Unsynchronized User Distribution with Non-compromising Adversaries. $z = 1 - (\frac{1}{2} + f_p(\ell))$, where $f_p(x) = \min(1/2, 1 - (1 - p)^x)$.

A derived δ lower bound for the non-compromising adversary is also a valid lower bound for a (partially) compromising adversary. For some edge cases (e.g., when ℓ is close to N and β is close to 0), due to some approximations employed in the compromising adversaries scenario, the non-compromising adversary lower bound is actually tighter than the compromising adversaries lower bound. Therefore, in Figure 10, while plotting the 3D graph for a partially compromising adversary scenario, we have used the maximum of the lower bounds on δ for compromising adversary and non-compromising adversary.

In each plot, the dark blue region indicates the possibility of obtaining strong anonymity. For any point (x, y) outside




Fig. 8. Synchronized User Distribution with Partially compromising Adversaries. Total protocol parties $K = 100$, number of compromised parties $c = 20$. $z = 1 - [1 - (\frac{c}{\ell}) / \binom{K}{\ell}] f_\beta(\ell)$ for $\ell \leq c$, $z = 1 - [1 - 1/\binom{K}{c}] f_\beta(c) - f_\beta(\ell - c)$ otherwise.




Fig. 10. Unsynchronized User Distribution with Partially compromising Adversaries. Total number of protocol parties $K = 100$, number of compromised parties $c = 20$. $z' = 1 - [1 - (\frac{c}{\ell}) / \binom{K}{\ell}] [\frac{1}{2} + f_p(\ell)]$ for $\ell \leq c$, $z' = (1 - [1 - 1/\binom{K}{c}] [1/2 + f_p(c)]) \times (1 - [1/2 + f_p(\ell - c)])$ otherwise. We set $z = \max(z', 1 - (1/2 + f_p(\ell)))$

those regions, strong anonymity is not possible. For example, as shown in Figure 7, for $\ell = 100$ the bandwidth overhead β has to be at least 0.01 to expect strong anonymity.

For the chosen c and K , the plots in Figures 7 and 8 are almost identical as the ℓ and β factors contribute more to anonymity than the compromised parties can affect it. If we instead compare Figure 9 with Figure 10, the effect of compromise is noticeable: the dark blue region in Figure 10 is much smaller than that in Figure 9. Also, we can see a steep wall in Figure 10 for $\ell \leq c = 20$, demonstrating that providing anonymity becomes difficult when $\ell < c$; however, for $\ell > c$, the effect of compromise is less noticeable.