

Access Granted, Privacy Lost: Formalizing & Quantifying the Hidden Anonymity Risks of Exclusive-Use Systems

Christopher Ellis
The Ohio State University
ellis.729@osu.edu

Zhiqiang Lin
The Ohio State University
zlin@cse.ohio-state.edu

Abstract

Exclusive-use systems emit binary interaction signals through core functions such as authentication, presence updates, and message submissions. Although sparse and encrypted, these signals reflect user-specific behavior and, when linked over time, can erode anonymity. Because each credential or device is uniquely tied to an individual, even minimal activity patterns can enable re-identification and behavioral inference, posing a hidden but persistent privacy risk. We present a formal framework that models this leakage by digitizing interaction outcomes into multidimensional binary signals and quantifying anonymity degradation using entropy-based Quantitative Information Flow (QIF), Bayes vulnerability, and indistinguishability games. To generalize the threat, we introduce a taxonomy spanning attacker capabilities, observation methods, and types of information leaked. After discovering these signals from network analysis of Microsoft Teams, we produce a simulation case study with varying user activity profiles, demonstrating that content-agnostic signals alone enable a passive adversary to achieve 54.7% Top-1 and 89.1% Top-3 re-identification accuracy in a 16-user pool, with mean entropy losses of approximately 1.2 bits (about 30% of the 4-bit anonymity space), and worst-case reductions exceeding 2.4 bits. Additional analyses of *WhatsApp* traffic and the *IDBleed* BLE relay attack highlight broader applicability. Our results show that binary observables long treated as benign can systematically compromise anonymity, establishing a cross-domain framework for formalizing, quantifying, and classifying privacy loss in exclusive-use systems. This framework further enables defenders to formally model exclusive-use systems and quantitatively evaluate the privacy impact of proposed mitigations using entropy and vulnerability-based metrics.

Keywords

User privacy, exclusive-use, network protocols, deanonymization, user profiling, tracking, formalization, privacy risks

1 Introduction

Modern systems often protect the contents of user interaction data, but overlook the simple fact that the occurrence of activity is observable in network communications. Many platforms emit traffic that can be interpreted as binary signals indicating interaction, without revealing message content or user identity. These include successful logins to a workplace VPN, bursts of encrypted traffic when




Figure 1: Digitization abstraction: observable traffic is matched to patterns, converted into binary signals, and accumulated into traces for inference.

user presence is detected, or “typing...” indicators shown during message composition. In isolation, these minimal signals may seem benign. However, in systems designed for *exclusive-use*, where each credential or device maps uniquely to an individual, even sparse activity patterns can reflect user-specific behavior over time and degrade anonymity. These observable traffic patterns, often inherent to system functionality, pose a serious but *often hidden risk to users*, as many do not offer a meaningful way to opt out of this side-channel information leak.

While encryption and statistical privacy techniques have substantially advanced the protection of message contents [18, 30] and aggregate data [19, 20], less work has examined how interaction-level signals expose users to long-term inference. In exclusive-use systems, where devices or credentials are uniquely tied to individuals, binary indicators such as login success, message submission, or typing status are often observable through encrypted traffic patterns [1, 37, 52]. These signals occur with high fidelity, reflect user actions, and may accumulate into persistent behavioral signatures.

Prior work on *IDBleed* [22] introduced the concept of exclusive-use and demonstrated how protocol-level success and failure signals can leak device affiliation. While this and other behavioral inference studies [40, 47, 50, 56] highlight concrete attacks, there remains little formalism for how anonymity degrades due to recurring binary signals. Existing frameworks typically focus on rich metadata, continuous traffic flows, or statistical aggregates. In contrast, exclusive-use systems emit semantically meaningful events that resist traditional mitigations and occur at discrete boundaries. These properties call for a distinct modeling and measurement approach to capture how anonymity degrades over time.

We address this gap by treating binary observables as a structured form of information leakage in exclusive-use systems. Our framework combines three components. First, we present a formal model that defines users, resources, and observability, and digitizes activity into multidimensional binary vectors. Second, we introduce a taxonomy of threats organized by attacker capability, observation method, and information leaked, which systematizes how inference

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

Proceedings on Privacy Enhancing Technologies 2026(3), 32–47

© 2026 Copyright held by the owner/author(s).

<https://doi.org/10.56553/popets-2026-0069>



arises from recurring observable signals. Third, we provide quantification methods that measure anonymity degradation through worst-case anonymity set reduction, entropy-based Quantitative Information Flow (QIF), Bayes vulnerability, and a behavioral indistinguishability game. Figure 1 illustrates this abstraction, in which observable traffic is matched to deterministic patterns, digitized into binary outcomes, and accumulated into behavioral traces suitable for adversarial inference. Together, these components enable defenders to model exclusive-use systems and quantitatively measure privacy loss before and after mitigation using entropy and vulnerability-based metrics.

During our investigation, we analyzed Microsoft Teams traffic in a controlled lab environment and discover deterministic exclusive-use signals. This inspires us to conduct a case study using simulated usage data based on previous studies to validate our approach. We simulate 16 users over four distinct activity profiles (deliberative, low engagement, high frequency, and intermittent) and evaluate both worst-case and probabilistic inference. Despite relying only on sparse binary signals from presence, typing, and message-sent traffic, a passive adversary achieves 54.7% Top-1 and 89.1% Top-3 re-identification accuracy in a 16-user pool. Mean entropy loss is approximately 1.2 bits (about 30% of the 4-bit anonymity space), with worst-case reductions exceeding 2.4 bits for certain profiles, demonstrating that even minimal signals can compromise anonymity in practice. We further show that the taxonomy extends naturally to additional contexts such as mobile messaging with *WhatsApp* and previously reported active relay attacks with BLE in *IDBleed*, underscoring the broader applicability of the framework.

Contributions. This paper makes the following contributions:

- **Formalization of Exclusive-Use Leakage (§3):** Building on prior definitions, we model how binary observables in exclusive-use systems form a structured information leakage side-channel.
- **Quantification Framework (§4):** We introduce complementary methods to quantitatively measure anonymity loss using worst case anonymity set reduction, entropy-based QIF, Bayes vulnerability, and a behavioral indistinguishability game.
- **Threat Taxonomy (§5):** We systematize attacker strategies across three dimensions (attacker capability, observation method, and information leaked) and map both simulated and real-world attacks into this space.
- **Case Study Evaluation (§6):** We apply our framework and show that sparse binary signals in Microsoft Teams traffic can support re-identification, yield significant entropy loss, and expose user behavior across pairwise and population-scale settings, and extending our analysis to additional applications.

2 Background

We provide essential background on system architectures, privacy frameworks, and attacker models relevant to the analysis of binary signal leakage. This includes exclusive-use systems, side-channel mechanisms, entropy-based inference models, and the role of taxonomies in structuring adversarial threat assessments.

Exclusive-Use Systems. An *exclusive-use system* is one in which each device, credential, or account is explicitly tied to a single user [22]. These systems include personal smartphones, enterprise

messaging accounts, smart home devices paired to an owner’s device, or virtual private network (VPN) clients linked to employee credentials. Exclusive-use systems differ from open-access platforms, such as shared computers or accounts (N:1 shared mappings), in that each system interaction can be attributed to a single entity or user. This relationship is typically one user to one resource (1:1) or one user to multiple exclusive resources (1:N). This property introduces persistent behavioral structure in interaction signals, as all activity originates from the same user or device. Prior work has shown that such systems can leak identity or affiliation even in the absence of unique identifiers, due to consistent and observable signal to functionality mappings, such as authentication [22].

Side-channels and Binary Observables. Side-channel attacks exploit observable system outputs or emissions that are not intended to carry semantic meaning (i.e., second order effects). These include low-level physical signals such as timing [43], power consumption [42], or electromagnetic radiation [27], as well as higher-level protocol metadata, structure, and patterns. More recently, attention has shifted to *behavioral side channels*, where useful signals are inherent in legitimate application and network-layer traffic [1]. These include patterns of activation, successful authentication, delivery receipts, or typing notifications. These signals are often observed as binary (success vs. failure) or categorical patterns, recorded as discrete time series. Even when temporally imprecise or anonymized with modern countermeasures such as MAC address randomization, these sequences can be digitized and abstracted into binary traces that retain the structure of user behavior over time.

Quantitative Information Flow (QIF). Quantitative Information Flow (QIF) provides a formal framework for reasoning about how much information an attacker learns from observing a system [14, 55]. Unlike one-time, ephemeral leakage, QIF captures the gradual reduction in uncertainty about a secret variable over time. Entropy-based metrics are commonly used, where a prior belief distribution over secrets is updated after observing outputs. The difference in entropy quantifies the information gained. Therefore, this model is well suited to scenarios where attackers observe repeated signals and seek to reduce the anonymity of the source in a population.

Indistinguishability Games. Indistinguishability games are commonly used in cryptography [31, 41] and differential privacy [19] to formalize whether two inputs can be reliably distinguished based on system outputs. In privacy analysis, they offer a test-based framework: an adversary is given two potential sources and a single observed trace and must guess which source produced it. This model supports probabilistic reasoning about how observable features, such as packet sequences or binary signal traces, enable attacker inference, even without full identity disclosure. Indistinguishability games are useful in evaluating the degree of separation between behavioral profiles under constrained observability.

Privacy Terminology. Foundational privacy terminology, such as that introduced by Pfitzmann and Hansen [48], defines key concepts including *anonymity*, *unlinkability*, and *unobservability*. They define anonymity as a target, or subject, being unidentifiable in a group of other targets, the anonymity set. Unlinkability ensures that multiple actions cannot be associated with the same user, while unobservability aims to hide that any action even occurred. These distinctions are critical in systems where content and identifiers

are protected, but the existence or timing of interactions remains visible. For example, a user may retain anonymity with respect to identity but become recognizable through repeated behavioral patterns, risking linkability of events. These definitions provide a consistent vocabulary for analyzing privacy degradation in systems where events traces, such as network traffic, are observable.

Attack Taxonomies and Threat Modeling. Attack taxonomies help systemize adversarial strategies by categorizing attackers along multiple dimensions that includes techniques and outcomes [11, 46, 61]. These frameworks are common in security and privacy research, offering insight into vulnerabilities and subsequent consequences based on attacker capabilities, visibility, and goals. In the context of behavioral leakage, taxonomies can capture how different adversaries exploit observable signals to infer identity, activity, or relationships with other entities of a system. They also support generalization by mapping concrete examples onto a structured framework, researchers can reason about threat coverage, mitigation scope, and transferability of attacks across systems.

3 Exclusive-Use Formal Model

We now formalize the structure of exclusive-use systems and the types of observable signals they produce. We accomplish this by modeling how interactions generate semantically meaningful binary signals, how these signals are digitized and observed, and how an adversary updates beliefs about user behavior over time. This framework forms the foundation for the quantification and evaluation presented in §4 and §6.

3.1 System Model

Let $\mathcal{U} = \{u_1, \dots, u_n\}$ be a set of users and $\mathcal{R} = \{r_1, \dots, r_m\}$ a set of resources such as devices, accounts, or applications, each controlled by a single user. We define an exclusive-use mapping:

$$\phi : \mathcal{U} \rightarrow 2^{\mathcal{R}}, \quad \text{where } \phi(u_i) \cap \phi(u_j) = \emptyset \text{ for all } i \neq j$$

This guarantees exclusivity with each resource tied to at most one user. The formulation supports both 1:1 mappings (one user, one resource) and 1:N mappings (one user, multiple resources), while explicitly excluding N:1 shared systems, which violate the attribution assumptions required for identity inference.

Time progresses in discrete steps:

$$t \in \{1, \dots, T\}$$

At each time step t , a user u may interact with their resource(s), producing a raw trace:

$$x_t^{(u)} \in \mathcal{X},$$

where $\mathcal{X}$ denotes the raw feature space (e.g., packet sizes, feedback states, metadata, timing sequences).

3.2 Digitized Semantic Signals

This model considers binary signals that represent discrete user interactions. These may originate as explicit binary observables (login success vs. failure packet types specified in a protocol), or may result from digitizing continuous general activity, such as a sequence of packets, into abstracted semantic indicators or signals. In either case, the result is a sequence of multidimensional binary observations that encode whether specific actions occurred at each

time step. We use *anonymity degradation* here to refer to reduced unlinkability or, equivalently, increased linkability between observed actions and the user identity.

Let there be N such interaction types, each represented by a digitization function:

$$D_j : \mathcal{X} \rightarrow \{0, 1\}, \quad j = 1, \dots, N$$

where $\mathcal{X}$ denotes the raw feature space (network traces, system logs), and each D_j extracts the presence or absence of a specific user interaction, such as a message sent indicator. The resulting digitized observation for user u at time t is:

$$\mathbf{o}_t^{(u)} = \left(D_1(x_t^{(u)}), D_2(x_t^{(u)}), \dots, D_N(x_t^{(u)}) \right) \in \{0, 1\}^N$$

This defines a multidimensional binary vector encoding the user's activity at time t .

The attacker's overall observable sequence is:

$$\mathcal{O}_T = \{\mathbf{o}_1, \mathbf{o}_2, \dots, \mathbf{o}_T\}, \quad \mathbf{o}_t \in \{0, 1\}^N$$

where each $\mathbf{o}_t$ is generated by some unknown user $u \in \mathcal{U}$. These binary vectors form the foundation for probabilistic inference, allowing attackers to reason about identity based on structured behavioral patterns over time.

Given a digitized behavioral trace $\mathcal{O}_T$, the attacker may extract a finite-dimensional feature representation $f(\mathcal{O}_T)$ capturing aggregate and structural properties such as event counts, session structure, and burst characteristics. Because we do not assume a generative probabilistic model of user behavior, we instantiate a distance-based likelihood surrogate. Specifically, for a candidate user u with reference feature vector θ_u , we define:

$$P(\mathcal{O}_T | u) \propto \exp \left(-\frac{d(f(\mathcal{O}_T), \theta_u)}{\tau} \right)$$

where $d(\cdot, \cdot)$ denotes Euclidean distance in normalized feature space, and τ is a scaling parameter ($\tau = 1$ in our experiments). Assuming a uniform prior over users, the attacker computes the posterior:

$$P(u | \mathcal{O}_T) = \frac{P(\mathcal{O}_T | u)}{\sum_{u'} P(\mathcal{O}_T | u')}$$

forming the basis for quantification in §4.

4 Privacy Quantification

To understand how exclusive-use signals reduce anonymity over time, we formalize metrics that capture attacker inference power under the posterior distribution defined in §3. Our framework includes: (i) a worst-case anonymity set reduction heuristic, (ii) entropy-based quantification from Quantitative Information Flow (QIF), (iii) Bayes vulnerability as an operational attacker-success metric, and (iv) a behavioral indistinguishability game. Together these provide complementary views of privacy degradation.

4.1 Worst-Case Anonymity Set Reduction

In some cases, attacker knowledge or system semantics allow immediate elimination of users whose behavior is incompatible with an observed signal via heuristics. Let $\mathcal{U} = \{u_1, \dots, u_n\}$ be the set of users, and $\mathbf{o}_t \in \{0, 1\}^N$ the observed signal at time t . The reduced anonymity set is:

$$\mathcal{U}_t = \{u_i \in \mathcal{U} : \mathbf{o}_t \text{ is consistent with user } u_i\}$$

Anonymity loss is then estimated as:

$$L_{\text{heuristic}}(\mathbf{o}_t) = \log_2 |\mathcal{U}| - \log_2 |\mathcal{U}_t|$$

This provides a conservative upper bound on inference power, assuming deterministic signal mapping and uniform priors. It captures the intuition that even a single binary observable may collapse the feasible anonymity set (e.g., a failed login excludes all but the user who actually failed).

4.2 Entropy-Based Quantification (QIF)

For probabilistic inference under uncertainty, we adopt the QIF framework. Let $S \in \mathcal{U}$ denote the unknown user. Before observing any signals, the attacker holds a prior:

$$P_0(S = u_i), \quad u_i \in \mathcal{U}$$

Unless otherwise noted, we assume uniform priors with baseline entropy:

$$H_0(S) = \log_2 n$$

representing an analytic baseline in which the attacker has no auxiliary information about likely identities. In practice, adversaries may hold biased priors informed by contextual knowledge such as user roles, expected working hours, historical activity levels, organizational structure, or network-level addresses/subnet allocations. Such information can further reduce initial uncertainty and amplify the effective leakage measured by our framework. Our use of a uniform prior therefore reflects a conservative starting point for analysis.

Given an observed trace $O_T = \{\mathbf{o}_1, \dots, \mathbf{o}_T\}$, the attacker computes the posterior distribution:

$$P(S = u_i | O_T) = \frac{P(O_T | u_i) P_0(S = u_i)}{\sum_{j=1}^n P(O_T | u_j) P_0(S = u_j)}$$

In our instantiation, the likelihood terms $P(O_T | u_i)$ are not assumed from a predefined probabilistic channel model. Instead, they are computed directly from observed behavioral traces by extracting feature representations $f(O_T)$ and measuring similarity to user-specific reference profiles, as defined in §3. Thus, the probabilities used in inference are derived from data rather than assumed a priori.

Privacy loss L is measured as Shannon entropy reduction in bits:

$$L_{\text{QIF}}(O_T) = H_0(S) - H(S | O_T)$$

with:

$$H(S | O_T) = - \sum_{i=1}^n P(S = u_i | O_T) \log_2 P(S = u_i | O_T)$$

Although the raw features (e.g., packet sizes, tuples) are not inherently binary, the digitization functions D_j ensure that only validated semantic events (e.g., presence updates, message sent) are mapped to $\{0, 1\}$. Thus, quantification operates over discrete, semantically meaningful signals rather than arbitrary traffic fluctuations.

As additional observations extend the trace from O_T to O_{T+1} , the posterior distribution is recomputed. The marginal impact of an additional signal can therefore be expressed as:

$$\Delta L_{T+1} = H(S | O_T) - H(S | O_{T+1}),$$

capturing how longitudinal exclusive-use signals progressively sharpen attacker belief.

Bayes Vulnerability. Entropy measures uncertainty reduction but does not directly quantify attacker success probability. To provide an operationally meaningful metric, we report Bayes vulnerability under the identity gain function:

$$V(O_T) = \max_{u \in \mathcal{U}} P(u | O_T).$$

Bayes vulnerability represents the success probability of an optimal one-shot attacker who guesses the most likely user given the observed trace. Unlike entropy, which measures expected uncertainty, Bayes vulnerability directly reflects re-identification confidence. This operational interpretation enables defenders to evaluate mitigation effectiveness by measuring how specific system changes reduce attacker success probability.

Illustrative Example. Consider a four-user system with uniform prior:

$$P(u_1) = P(u_2) = P(u_3) = P(u_4) = 0.25$$

After observing trace O_T , suppose the posterior becomes:

$$P(u_1 | O_T) = 0.70, \quad P(u_2 | O_T) = 0.15$$

$$P(u_3 | O_T) = 0.10, \quad P(u_4 | O_T) = 0.05$$

The prior entropy is:

$$H_0 = \log_2 4 = 2 \text{ bits}$$

The residual entropy is:

$$H(S | O_T) = - \sum_{i=1}^4 P(u_i | O_T) \log_2 P(u_i | O_T) \approx 1.32 \text{ bits}$$

Entropy loss is:

$$\Delta H(O_T) = 2 - 1.32 = 0.68 \text{ bits}$$

Bayes vulnerability is:

$$V(O_T) = \max_{u \in \mathcal{U}} P(u | O_T) = 0.70$$

indicating that an optimal attacker succeeds with 70% probability.

4.3 Indistinguishability Games

We define a binary decision game in which an attacker is given a signal sequence O_T generated by one of two known user profiles, u_0 and u_1 . One of the profiles is selected at random (i.e., $b \in \{0, 1\}$), and $O_T \sim \pi_b$. The attacker must guess the value of b by selecting the more likely source:

$$\text{Guess} = \arg \max_{i \in \{0,1\}} P(O_T | u_i)$$

The indistinguishability advantage is defined as:

$$\text{Adv} = 2 \cdot \left| \Pr[\text{Guess} = b] - \frac{1}{2} \right|$$

This metric quantifies how distinguishable two user profiles are based on structured observation. Unlike entropy and Bayes vulnerability, which measure population-wide uncertainty and re-identification confidence, this game isolates pairwise behavioral separability between specific users.

4.4 Inference Methodology

Given an observed digitized trace O_T , the attacker extracts a feature vector $f(O_T)$ capturing aggregate and structural behavioral properties (e.g., event counts, burst structure, and session gaps). For each candidate user u , a reference feature vector θ_u is constructed from prior traces. Likelihoods $P(O_T | u)$ are computed using the distance-based surrogate defined in §3, which converts similarity in feature space into a probabilistic score. Bayes' rule then yields the posterior distribution $P(u | O_T)$, from which entropy loss and Bayes vulnerability are computed. This pipeline operationalizes how observed exclusive-use signals translate into quantified anonymity degradation.

5 Threat Taxonomy

Exclusive-use systems inherently bind observable signals to specific users, making even seemingly benign binary events, such as authentication or access attempts, presence states, or message notifications, semantically meaningful. These signals, though minimal in isolation, accumulate into behavioral traces that can compromise anonymity over time. To systematically characterize adversarial strategies enabled by such observables, we build on the privacy terminology of Pfitzmann and Hansen [48], extending concepts such as anonymity, unlinkability, and unobservability to settings where deterministic, identity-linked feedback supports linkability and profiling.

Our taxonomy organizes attacks along three dimensions: *attacker capability* (who and when), *observation method* (how), and *information leaked* (what). Each path across these dimensions corresponds to a concrete strategy for degrading anonymity in exclusive-use systems. We present the taxonomy in Table 1 along with more detailed descriptions of the paths for the remainder of this section as a tool for researchers to apply and model real-world scenarios.

Table 1 grounds the taxonomy in real-world systems, showing paths across diverse platforms and adversarial vantage points. The examples span from passive sniffing of messaging traffic, to probing in IoT systems, to privileged log access in enterprise environments. This breadth highlights that the taxonomy is not limited to theoretical constructs: it captures adversarial strategies already documented across domains. By positioning known attacks within the same structure, the table demonstrates both the generality of the taxonomy and its utility for analyzing new threats.

5.1 Attacker Capability

Observation Time. The attacker may operate in a short-lived setting with only a single opportunity to observe a signal (*Singular*), or may collect signals across multiple sessions to support re-identification and pattern analysis (*Longitudinal*). Longitudinal visibility amplifies inference by allowing the attacker to correlate behavior over time.

Access Level. Attackers may be limited to externally observable signals and public interfaces (*Unprivileged*), or may access internal logs, diagnostics, or telemetry unavailable to standard users (*Privileged*).

Locale. Attackers may be nearby in proximity, or embedded in the same physical or network environment as the user (*Local*), enabling

high-fidelity signal capture and timing inference. Alternatively, operating from a distance (*Remote*), relying on signals exposed through public-facing protocols or interfaces.

Scope. Some attackers focus on a single platform or application (*Mono-system*), while others correlate behavior across different platforms or services (*Cross-system*). Cross-system attackers pose a powerful threat by matching behavioral fingerprints across otherwise isolated contexts.

5.2 Observation Method

Passive. Passive attackers observe signals without interacting with the system or triggering responses. This includes capturing data from shared or open communication channels (*Sniffing*), monitoring developer-facing analytics that expose state transitions (*Exposed Metrics*), correlating observable event timing such as session joins or authentication attempts (*Timing*), or accessing application or system logs that capture user behavior (*Logs*), or building unique or identifiable profiles of users or devices from observing legitimate traffic patterns and behavior (*Fingerprinting*). These methods exploit feedback already present in the system without generating new activity. In practice, passive sniffing may occur at shared subnets, enterprise gateways, or institutional monitoring points where encrypted traffic metadata is routinely observable.

Active. Active attackers interact with systems to create responses. This may include forwarding data between users or devices (*Relay*), re-sending previously observed packets to manipulate system state (*Replay*), sending crafted queries to elicit success/failure or other binary responses (*Probing*). Active observation expands visibility but increases the risk of detection.

5.3 Information Leaked

User Inference. The attacker may resolve a user's identity through behavioral matching or direct confirmation (*Identity Resolution*), or determine whether a user belongs to a specific group or access-controlled context (*Group Membership*). These inferences often require only minimal visibility and are amplified in exclusive-use systems.

Relationship Inference. Binary interaction signals may reveal 1:1 entity relationships (*Pairwise Linking*) or support reconstruction of larger social or organizational structures based on coordinated patterns of activity (*Association Mapping*). These inferences undermine unlinkability and enable network-level profiling.

Behavioral Inference. Attackers can extract rich behavioral information from repeated signals. They may observe when a user is online or active (*Presence Detection*), infer the structure and frequency of activity sessions (*Session Attributes*), profile the user's engagement style based on signal types and frequency (*Role Characterization*), or construct a consistent signature of user behavior over time (*Behavioral Fingerprint*). Exclusive-use systems heighten the impact of behavioral inference as every signal inherently corresponds to a single user or entity.

5.4 Applying the Taxonomy

The taxonomy is best applied by tracing a path across its three dimensions: start with the attacker's *capability*, then specify the

Attacker Capability	
Observation Time	
— Singular	Inferring association from one-time authentication [22]
— Longitudinal	Deanonymizing users in a messaging platform based on recurring activity signals (this paper)
Access Level	
— Unprivileged	Observing typing and presence indicators via standard interfaces in encrypted platforms (this paper)
— Privileged	Inferring user activity via admin-accessible presence APIs or logs [45]
Locale	
— Remote	Extracting website user activity from encrypted traffic traces [37]
— Local	Monitoring wired/wireless network traffic to infer user activity (this paper)
Scope	
— Mono-system	Performing inference using metadata within a single encrypted messaging platform (this paper)
— Cross-system	Modeling user behavior across multiple subsystems by combining feature vectors [56]
Observation Method	
Passive	
— Exposed Metrics	Polling presence or status endpoints exposed by APIs [29]
— Fingerprinting	Differentiating users or traffic sources based network traffic structure [37]
— Logs	Inferring user interest based on website logs [45]
— Sniffing	Capturing encrypted wired/wireless network traffic to infer user activity [1]
— Timing	Inferring spatial user behavior from observed packet timing offsets [1]
Active	
— Relay	Replaying Bluetooth pairing sequences to emulate physical proximity [22]
— Replay	Triggering connectivity by retransmitting previously captured packets [65]
— Probing	Detecting valid accounts by querying servers [28]
Information Leaked	
User Inference	
— Identity Resolution	Matching behavioral patterns across sessions to deanonymize users (this paper)
— Group Membership	Inferring group membership in aggregate location datasets [50]
Relationship Inference	
— Pairwise Linking	Inferring 1:1 interaction from presence/message interleaving [40]
— Association Mapping	Reconstructing social relationships through activity timelines or use [22]
Behavioral Inference	
— Presence Detection	Monitoring online availability signals over time (this paper)
— Session Attributes	Extracting session length and burst gaps from encrypted metadata (this paper)
— Role Characterization	Characterizing user role and/or activity based on multiple dimensions (this paper)
— Behavioral Fingerprint	Re-identifying users through recurring binary signal structure (this paper)

Table 1: Taxonomy of attacker strategies in exclusive-use systems. The table organizes exemplar attacks by *attacker capability*, *observation method*, and *information leaked*, showing how concrete examples map to each node in the taxonomy.

observation method, and finally identify the *information leaked*. Each complete path describes a concrete strategy. For instance, an *Unprivileged, Local, Longitudinal* adversary using *Passive Sniffing* may derive *Behavioral Fingerprints* from presence and typing bursts in messaging traffic. In §6, we map the taxonomy onto both our Microsoft Teams case study as well as brief additional examples, including *WhatsApp* event indicators and the *IDBleed* active BLE relay attack, to highlight generality.

6 Case Study: Microsoft Teams

We apply our exclusive-use formalization (§3), privacy quantification (§4), and taxonomy (§5) to a simulated case study grounded in behavioral patterns observed in real-world messaging platforms [57]. Observable interaction signals, such as presence activity, typing indicators, and message sends, are digitized into binary traces that reflect user-specific activity over time. We focus on Microsoft Teams as a widely deployed exclusive-use collaboration tool in academic

and corporate environments, where these signals are directly tied to individual identities. While these binary signals are relatively sparse over time, we demonstrate that their structured recurrence and semantic meaning degrades anonymity and enables inference.

We selected Microsoft Teams because, during exploratory traffic analysis, we noticed that certain packet patterns aligned with user actions such as becoming active, typing, or sending a message. This was a striking extension of the exclusive-use leakage: not just low-level protocols, but mainstream collaboration platforms were emitting binary signals that mapped directly to individual behavior. Teams was also a natural choice because of its prevalence in enterprise and academic settings, where accounts are closely tied to specific individuals. Although ethical and practical constraints prevent us from collecting real user traffic at scale, these initial observations in our controlled lab environment motivated our design of simulated user profiles grounded in realistic user activity that capture the same semantic structure without requiring sensitive data.

The evaluation proceeds in four stages. First, we produce a high-fidelity user activity simulation based on four behavioral profiles. Next, we apply a series of indistinguishability games across user pairings and populations, quantifying entropy loss as a metric of anonymity degradation. The analysis is then expanded to full pool-based inference scenarios with sixteen users. Finally, we evaluate re-identification directly via posterior rank across multi-day traces, further supporting the practical impact of our framework. Throughout, we report means and dispersion (standard deviation), reference ΔH in bits, and provide Bayes vulnerability as a metric of operational attacker success, as defined in §4.

6.1 Threat Model & Binary Signals

We consider a passive adversary capable of observing encrypted Microsoft Teams traffic at the network layer, such as access to the same Wi-Fi network or within an organizational infrastructure. The attacker cannot decrypt contents but can observe packet sizes and timings to infer application-layer interactions. The attacker capability is *Longitudinal, Unprivileged, Local, Mono-system*, with no endpoint instrumentation, key material, or API access with the observation method as *Passive Sniffing*. The information leaked is user and behavioral inference, including *Identity Resolution, Presence Detection, Session Attributes, Role Characterization*, and *Behavioral Fingerprint*.

Concretely, such observation may arise from several realistic vantage points. In small office or academic settings, an attacker may be positioned on the same local subnet (shared Wi-Fi or Ethernet segment) and passively capture encrypted traffic. In enterprise environments, similar visibility may exist at organizational gateways, firewall appliances, SD-WAN edges, or network monitoring systems that log packet metadata without decrypting contents. Because Teams traffic is routed through centralized infrastructure, these positions naturally enable longitudinal trace collection over days or weeks without endpoint compromise or privileged API access.

Digitized Observables from Packet Structure & Patterns. We analyzed network traffic transmitted from the Teams web client during typical feature use and identified deterministic traffic patterns corresponding to presence changes, typing indicators, and message submissions. This was conducted in our controlled lab subnet, using the Google Chrome web browser and network traffic inspection tool and Wireshark, running on Ubuntu 22.04. Specifically, and as shown as in Table 2:

- **Presence:** A two-packet sequence of sizes 184 and 250 bytes is emitted when a user becomes active after an inactive state (No. [21, 22]).
- **Typing:** A recurring tuple begins with a variable-sized packet X_i , consistent per message, followed by a fixed 205-byte packet and an 89-byte packet (No. [45, 46, 57], [102, 103, 116], and new message X_{i+1} [217, 218, 230]).
- **Message Sent:** Each send indicator terminates the current message using X_i (No. [125, 126]).

Each digitized dimension relies on strict pattern-matching criteria. For example, the *Typing* indicator is emitted only when the full tuple $\langle X, 205, 89 \rangle$ is observed in temporal sequence. Deviating sequences are excluded to avoid false positives to ensure that only

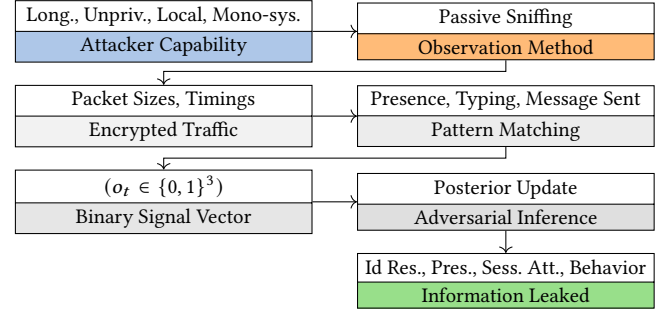


Figure 2: Taxonomy-augmented digitization pipeline for Microsoft Teams.

active typing windows are captured. The *Presence* and *Msg Sent* dimensions use similarly restrictive patterns. The digitization to binary signals of these dimensions are illustrated in Figure 2. Although the raw features (e.g., packet sizes, tuples) are not inherently binary, the digitization functions D_j map only validated semantic events (presence, typing, message sent) to $\{0, 1\}$, ensuring our quantification targets discrete, semantic signals rather than arbitrary traffic fluctuations.

No.	Source IP	Destination IP	Pkt. Size	Inference
21	<User IP>	<MS Teams IP A>	184	Presence
22	<User IP>	<MS Teams IP A>	250	Presence
...				
45	<User IP>	<MS Teams IP B>	2937	Typing
46	<User IP>	<MS Teams IP B>	205	Typing
57	<User IP>	<MS Teams IP B>	89	Typing
...				
102	<User IP>	<MS Teams IP B>	2937	Typing
103	<User IP>	<MS Teams IP B>	205	Typing
116	<User IP>	<MS Teams IP B>	89	Typing
...				
125	<User IP>	<MS Teams IP B>	2937	Msg Sent
126	<User IP>	<MS Teams IP B>	928	Msg Sent
...				
217	<User IP>	<MS Teams IP B>	3125	Typing
218	<User IP>	<MS Teams IP B>	205	Typing
230	<User IP>	<MS Teams IP B>	89	Typing

Table 2: Size patterns of real-world Teams user activity network packet traces and illustrating X_{i+1} at packet 217. IP addresses redacted with representative labels for a single user and consistent MS Teams IP destinations.

6.2 Simulation Setup

We simulate a workplace environment consisting of sixteen ($N = 16$) users across four behavioral profiles: *Deliberative* (Alice), *Low Engagement* (Bob), *High Frequency* (Charlie), and *Intermittent* (Diane). Each user generates an 8-hour activity trace to simulate an average workday, including presence, typing, and message sent

events. Event frequencies are anchored to published messaging analytics (approximately 300 messages per user per week [57]), producing empirically grounded activity distributions. Session bursts, inter-event timing variability, and user-specific activity rates are modeled to reflect common workplace communication patterns and application behavior rather than uniformly random or synthetic traffic. Presence events are injected after inactivity, typing events may be abandoned, and messaging sequences reflect real-time interaction bursts. We simulate ten runs per user and profile to capture behavioral variability, dropping the top and bottom outliers during calculations.

We select $N = 16$ to provide a meaningful anonymity baseline while maintaining experimental clarity. With $N = 16$, the initial uncertainty is $H_0 = \log_2 16 = 4$ bits, allowing us to quantify substantial entropy reductions in an interpretable setting. This population size reflects a realistic set of candidate users observable within a single workspace, such as a small office, lab, or department on a larger enterprise subnet, while the framework itself generalizes to larger populations without modification. Figure 9 provided in the Appendix presents a rasterized timeline for all users, visually highlighting the structure and diversity of the activity traces to support the simulation realism. This setup keeps the semantics of exclusive-use intact while allowing controlled variation in session structure, burstiness, and idle intervals.

6.3 Indistinguishability Games (4 Profiles)

We evaluate directional indistinguishability games across four behavioral profiles: Alice (A), Bob (B), Charlie (C), and Diane (D). For each pair, we report both directions (e.g., A^* vs B and B^* vs A), where the starred user denotes the true source of the observed trace. The y-axis in Figure 3 shows the posterior probability assigned to each candidate, with ΔH annotated in bits. In the two-user setting, Bayes vulnerability equals the maximum posterior, thus, when the true user (*) receives highest probability, it matches the reported posterior.

Cross-profile pairs are consistently separable. Alice vs Bob yields posteriors of 0.85 ± 0.08 and 0.87 ± 0.05 , with ΔH between 0.42 and 0.45 bits. Bob vs Diane and Alice vs Diane show the strongest separation, with posterior probabilities above 0.97 and entropy loss up to 0.90 bits. Charlie is likewise well distinguished from Bob and Diane (ΔH between 0.58 and 0.63 bits). Alice vs Charlie exhibits the weakest separation (posterior ≈ 0.71 , ΔH between 0.16 and 0.19 bits), reflecting closer behavioral similarity. Directional differences in ΔH indicate that some profiles produce more distinctive traces than others despite equalized message volume.

6.4 Expanded Evaluation ($N = 16$)

To generalize the pairwise evaluation, we simulate four users per profile ($N = 16$) and compute all 120 directional games. We present a truncated view showing the top, middle, and bottom three pairs by posterior confidence assigned to the true user in Figure 4, and provide the full set in Figure 8.

High-confidence cross-profile pairs approach near-deterministic inference. For example, B_3^* vs D_3 and its reverse both exceed 0.997 posterior with ΔH above 0.97 bits. Similar B_3 - D_1 combinations




Figure 3: Posterior probabilities in directional indistinguishability games between user profiles. ΔH denotes entropy loss in bits.




Figure 4: Posterior probability assigned to the true user in selected pairwise games ($N = 16$).

consistently exceed 0.99 posterior, indicating almost complete separability. Mid-ranked pairs remain clearly distinguishable but less extreme. For instance, B_2^* vs C_1 yields $P(\text{true} | O) = 0.89$ with $\Delta H = 0.55$ bits, and B_1^* vs B_4 yields $P = 0.90$ with $\Delta H = 0.54$ bits. In contrast, some same-profile pairs collapse toward chance-level inference. Examples such as C_1^* vs C_3 and C_1^* vs C_2 produce posteriors near 0.49 and $\Delta H < 0.03$ bits. These results confirm that anonymity erosion scales with behavioral diversity: cross-profile traces yield rapid entropy reduction, whereas similar behavioral profiles remain difficult to distinguish.

To validate that the simulated users exhibit structured behavioral variation, we perform PCA on the normalized feature vectors derived from digitized binary traces. As shown in Figure 7, the embedding reveals separation aligned with the four behavioral profiles. The first principal component explains 47.6% of the variance and the second explains 18.9%, indicating that behavioral differences span multiple dimensions rather than collapsing onto a single dominant feature. This supports that the simulation produces distinct activity patterns consistent with the attacker's multidimensional observation model.




Figure 5: Mean entropy loss (ΔH) and posterior probability per user in a 16-user pool.




Figure 6: Posterior probability and rank of the true user in leave-one-out re-identification.




Figure 7: Principal Component Analysis (PCA) of behavioral features extracted from digitized traces.

6.5 Pool Re-identification (N=16)

We evaluate population-scale re-identification with an anonymity set of $N = 16$ with four users per profile. For each target trace, the attacker computes a posterior over all candidates using reference traces from other runs. Figure 5 reports mean entropy loss ΔH and Bayes vulnerability V per user. Across users, observed binary signals reduce uncertainty by approximately 0.68 to 2.44 bits (mean 1.2 bits or roughly 30%) out of the 4-bit anonymity space.

Deliberative and high-frequency users (A and C) exhibit moderate entropy loss (≈ 0.78 -1.00 bits), while low-engagement and intermittent users (B and D) experience substantially larger reductions. The strongest case, D_2 , loses 2.44 bits on average, shrinking the effective anonymity space from 16 users to fewer than 3. Bayes vulnerability correspondingly increases from the uniform prior of $1/16 = 6.25\%$ to between 19% and 71%, depending on behavioral distinctiveness. Several D_i and B_i users exceed 0.45, indicating near one-shot identification from behavior alone.

We also report re-identification performance in terms of the true user's rank under the posterior. Figure 6 plots, for each evaluated trace, the posterior probability assigned to the true identity versus its rank among all $N = 16$ candidates. Top-1 accuracy is 54.7% and top-3 accuracy reaches 89.1%, with mean rank 2.23 and median rank 1. The mean posterior assigned to the true identity is 0.275 (uniform prior $= 1/16 = 0.0625$), with maxima approaching 0.832. On average, this represents an increase of approximately 21 percentage points over random guessing (from 6.25% to 27.5%), with the most distinctive users experiencing gains exceeding 64 percentage points. Together, these results show that even coarse binary interaction signals substantially reduce anonymity and enable high-confidence identification within a 16-user population.

6.6 Teams Taxonomy Mapping

Concrete taxonomy paths illustrate the Teams threat model. The adversary's *Capability* is *Longitudinal, Unprivileged, Local, Mono-system*: they are simply another observer on the network, without special credentials, but positioned to collect traces over extended periods within the same platform. The *Observation Method* is *Passive Sniffing*, relying only on encrypted packet sizes and timings. The resulting *Information Leaked* spans *Behavioral Inference* with presence detection, session attributes, and behavioral fingerprints, as well as *User Inference* via identity resolution. Positioned this way, the taxonomy makes explicit how limited observability and minimal signals accumulate into re-identification power in exclusive-use environments such as Teams.

6.7 Additional Applications of the Taxonomy

We briefly illustrate two additional applications, *WhatsApp* message traffic and the *IDBleed* BLE active relay attack, to demonstrate the generalization of the taxonomy to other domains.

Taxonomy: *WhatsApp*. In controlled tests, we observed that *WhatsApp* produces distinctive encrypted traffic bursts, with packet size and sequencing patterns reflecting typing, individual chat focus changes, and message send events from a unique IP to a *WhatsApp* application IP. Specifically, we notice a packet of size 134 bytes when a user begins typing to another single user (136 bytes for groups) and another packet of equal size when they stop. Similarly, packets of sizes 211-213 bytes when changing chat focus, and a packet size of 101 bytes following a larger packet containing the message data, after a sent message ($\langle \langle msg \rangle, 101 \rangle$). Although content is protected, these repeated structures can be digitized into binary observables in the same manner as Teams.

Mapping through our taxonomy, the attacker has *Longitudinal, Unprivileged, Local, Mono-system* capability, applies *Passive Sniffing*, and infers *Behavioral Fingerprints* and *Identity Resolution*. Even a




Figure 8: Pairwise indistinguishability results for all 120 user pairs ($N = 16$), with the clear separation between users with different profiles reflecting varying degrees of anonymity degradation.

single conversation partner can therefore leak recognizable activity patterns.

Taxonomy: IDBleed. The *IDBleed* work demonstrated that Bluetooth Low Energy (BLE) pairing protocols between trusted, exclusive-use devices expose binary success/failure outcomes that can be exploited through an active relay attack to track users [22]. An attacker relays packets between trusted devices to probe connection attempts, learning whether authentication succeeded, subsequently deanonymizing the device and user.

Applying our taxonomy, this corresponds to *Singular, Unprivileged, Remote, Mono-system* capability, using an *Active Relay* method, leaking both *Group Membership* (is the target associated with this device?) and *Identity Resolution* (linking the user to a randomized BLE MAC address after observing the authentication success/failure response). This case complements our *Teams* and *WhatsApp* evaluations by showing that exclusive-use leakage extends beyond messaging applications into lower-level protocols, reinforcing the generality of our formal model and taxonomy.

7 Mitigation Strategies

Defending against anonymity degradation in exclusive-use systems requires mitigations to constrain adversarial capabilities, limit observability, or reduce information leakage, corresponding with our taxonomy dimensions (§5).

Limiting Attacker Capability. Restricting access to internal diagnostics, telemetry feeds, and interface state is a critical mitigation against privileged adversaries. Enforcing strict access controls, minimizing debug output in production environments, and auditing API exposure can significantly limit internal signal leakage [29, 46]. For example, isolating data available from user-accessible domains can prevent attackers from inferring presence or activity through privileged vantage points. Systems should adopt the principle of

least privilege not only for users but also for internal services, especially in mobile and embedded environments where such channels are often exposed by default.

To defend against local attackers, such as co-resident adversaries on shared networks or physically proximate sniffers, systems must reduce the fidelity of externally observable signals. Techniques such as traffic shaping [16, 64] and cover traffic [33] can obscure timing relationships and suppress identifiable bursts. Randomized transmission intervals and padding mechanisms can degrade the precision of passive observers [16, 54]. These defenses are particularly relevant in corporate or campus environments where unprivileged local observers can easily monitor broadcast or multicast channels.

Cross-system correlation is a powerful capability that allows attackers to match behavioral fingerprints across otherwise siloed services. Limiting the uniqueness of interaction patterns, especially at the binary signal level, is essential. Techniques such as random traffic injection to suppress behavior templates or fingerprinting [12] and deliberate introduction of ambiguity into interaction streams [22] can reduce inter-context linkability. Additionally, systems can randomize feedback mechanisms such as presence indicators to make them less temporally consistent across services. Implementing asynchronous interactions or batching data can help break cross-system timing [7].

To resist longitudinal inference, systems should adopt ephemeral identifiers, limit persistent behavioral exposure, and rotate observable patterns where feasible [4, 22–24, 32, 36, 38, 54]. Session-based obfuscation, such as time-bounded access tokens, location nonces, or short-lived activity logs, can help constrain the attacker’s observational window [21]. Periodically resetting feedback mechanisms or introducing noise over long time spans ensures that consistent patterns are less effective in inferring reliable identifiers due to traffic analysis.

Limiting Observation Methods. Because passive attackers exploit signals already exposed by the system, defenses must focus

on minimizing external observability and limiting the precision of leaked information. To counter sniffing and timing-based inference, traffic-level defenses such as padding, shaping, and randomized delays can obscure binary signal boundaries [16, 63]. For example, shaping login or message delivery responses to occur at fixed intervals regardless of true event timing prevents adversaries from correlating presence or activity with high fidelity in large user systems. To mitigate leakage from logs or exposed metrics, systems should sanitize diagnostic outputs, enforce strict audit controls, and disable unnecessary telemetry in production environments [26, 49]. Systems should treat all externally visible state changes as potentially sensitive and enforce minimization-by-default design or configuration.

To mitigate active relay and replay attacks, systems should enforce nonce-based freshness checks and rate-limiting policies that prevent reuse of prior messages [9, 39]. Mitigating probing requires input validation, rate-limiting, conservative error handling, and proper access controls to avoid revealing semantic characteristics or functionality. For fingerprinting, systems should randomize protocol-level identifiers, response sizes and timing, or feature orderings, making it harder for attackers to distinguish users or devices based on interactions [22, 63]. System designers should assume that even low-rate interactions can leak information and build defenses to prevent unintentional unique metadata or telemetry.

Reducing Information Leakage. Exclusive-use systems pose a high risk for user inference, as every signal is attributable to a single individual. To mitigate identity resolution and group membership attacks, systems can reduce the uniqueness of individual behavior patterns and limit metadata exposure. One practical strategy is to introduce behavioral ambiguity through randomized or obfuscated signal mechanisms [3]. Another approach is to decouple signals from stable identifiers by using rotating session identifiers or unlinkable authentication tokens [8]. These techniques prevent adversaries from linking binary events back to a persistent identity, reducing both direct and statistical re-identification attacks.

Binary signals that co-occur across users or sessions can expose interpersonal relationships or support broader association mapping. To mitigate pairwise linking, systems should randomize or suppress observable timing patterns that reflect direct exchanges [64]. Buffered delivery, message batching, and varying messaging intervals help desynchronize communication events that would otherwise reveal social ties [7, 15]. For larger-scale association mapping, designs can adopt indirection techniques, such as server-mediated batching [59], that blur the group topology reflected in user activity. These defenses reduce the ability to observe interaction sequences and reconstruct relational graphs.

Behavioral inference is particularly difficult to defend against in exclusive-use systems due to the semantic weight of each signal. To mitigate presence detection and reduce inference accuracy, systems can strategically introduce uncertainty into observable signals by coarsening timestamps, grouping signals into activity windows, or injecting synthetic traffic [22]. Role characterization and user fingerprinting can be mitigated by reducing behavioral distinguishability and anonymity loss, for example through padding, randomized interaction sequences, or cover traffic that make users less separable in their usage profiles [33]. Importantly, systems

should limit deterministic feedback such as typing indicators or read receipts, as these signals introduce predictable structures that can persist across sessions.

8 Discussion

Using the Framework for Mitigation Evaluation. The framework is designed not only to characterize exclusive-use leakage, but to measure the impact of mitigations. Because interaction signals are explicitly digitized and inference is quantified through entropy reduction and Bayes vulnerability, defenders can evaluate a system before and after a design change and directly compare resulting anonymity loss. A mitigation that reduces ΔH , lowers vulnerability, or shifts true-rank outcomes toward uniform behavior provides measurable evidence of improved privacy. The taxonomy further clarifies which attacker capabilities and observation methods are being constrained. Together, this enables a quantifiable and actionable workflow for identifying and reducing anonymity risk in exclusive-use systems.

Applicability Beyond Case Studies. Exclusive-use leakage is not limited to messaging platforms. Our framework naturally extends to enterprise VPNs, cloud APIs, smart systems, physical access controls, and other systems where observable user or device activity such as authentication, bi-directional data flows, and system state serve as binary signals. These observable signals, while essential for functionality, can expose identity or group affiliation when linked over time. Our framework provides a formal toolset to quantitatively reason about risks and model threats across these exclusive-use environments.

Assumptions and Limitations. Our case study evaluation models passive adversaries and assumes accurate digitization of semantic signals from encrypted traffic. While this assumption is grounded in prior work and validated through controlled simulations, real-world deployments, particularly with mitigations previously discussed, may introduce variability that complicates signal extraction. In addition, the behavioral traces used in our analysis are simulated, based on published analytics of messaging usage rather than real-world, empirical traffic captures. While this approach limits our ability to generalize findings to more diverse populations or operational environments, it reassuringly preserves user privacy and supports controlled evaluation to demonstrate the significant degradation of anonymity from binary signals.

Timing-Aware Inference vs. Volume-Only Baseline. To evaluate the importance of temporal structure, we compare our timing-aware feature model against a naive baseline that relies only on aggregate event counts (presence, typing, and message totals) without session or burst structure. Under true-rank re-identification, the timing-aware model achieves 54.7% Top-1 accuracy and 89.1% Top-3 accuracy, with mean rank 2.23 and median 1. In contrast, the volume-only baseline achieves 44.5% Top-1 accuracy and 72.7% Top-3 accuracy, with mean rank 2.59 and median 2. Overall, timing-aware inference improves Top-1 accuracy by 10.2 percentage points and Top-3 accuracy by 16.4 percentage points, confirming that re-identification is driven not only by the amount of binary signals, but by when and how they occur over time.

Cross-System Correlation and Advanced Adversaries. The general threat model can be extended to adversaries who correlate activity across systems. For example, an attacker observing VPN authentication timestamps might match them with other binary signals from messaging traffic to identify remote employees, allowing for a reduction in the anonymity set. Even if each system preserves anonymity with their own mechanisms, evaluating the intersection of compiled data can significantly reduce attacker uncertainty. This greater threat from advanced adversaries with broader access underscores the need for privacy analysis that covers multiple exclusive-use systems making up a larger ecosystem. Advanced adversaries may also attempt to utilize deeper statistical analysis or machine learning. While such approaches can scale to large deployments, they often require structured inputs or labeled training data. Our formalization offers a principled foundation for defining inference goals and measuring anonymity degradation in ways that complement or guide these techniques. Rather than replacing formal analysis, machine learning may benefit from the structure and interpretability our framework provides.

Imposed Consent. Our findings expose a deeper design flaw in exclusive-use systems: *they force a model of imposed consent*. Because these binary signals, and network traffic in general, are effectively leaked emissions from essential system operations, users cannot meaningfully opt out. Every interaction contributes to a behavioral fingerprint, often without awareness or recourse, resulting in involuntary identity exposure. We argue that this systemic leakage raises significant tension with data protection principles such as those outlined in the General Data Protection Regulation (GDPR) [25], which emphasize informed consent and data minimization. Exclusive-use systems silently expose user and device interaction data at and below the application layer, typically lacking user disclosure or control beyond toggling superficial settings like read receipts or typing indicators. While such systems may not store explicit identities, the persistent observability of interaction patterns introduces a meaningful privacy risk that warrants scrutiny and reconsideration under evolving interpretations of data protection law.

Impact and Consequences. Our findings reveal that digitized binary signals in exclusive-use systems, even when anonymized and encrypted, can carry behavioral fingerprints strong enough to compromise anonymity. The entropy losses observed in our pairwise and population-scale experiments from our Microsoft Teams case study translate into meaningful privacy erosion: a reduction in attacker uncertainty, an increase in confidence, and the potential for profile linkability across time or systems.

These attacks do not require content inspection or explicit user identifiers. Instead, they operate entirely on structural patterns derived from interaction signals. In practical terms, an adversary monitoring Teams traffic could probabilistically associate typing bursts with individual users, narrow candidate sets across messages, or link session behavior across applications. The risk compounds over time, and across platforms, making exclusive-use systems uniquely vulnerable to these hidden anonymity threats. Moreover, in operational environments a local observer can often correlate link- or network-layer identifiers (e.g., MAC/DHCP/IP bindings available to infrastructure operators) with observed behavioral

fingerprints, enabling eventual deanonymization of the individual behind a device even when only encrypted traffic is visible.

Future Work. Several directions warrant further research. One is the development of leakage scoring metrics to quantify system vulnerability under varying attacker capabilities, enabling comparative evaluations across architectures. Another is the formal modeling of plausible deniability and uncertainty windows, capturing conditions under which users retain ambiguity despite observable interaction signals. In addition, future work could apply this framework to real-world datasets under privacy-preserving conditions, enabling empirical validation of behavioral leakage in deployed systems. Cross-system evaluation is another promising direction, examining how exclusive-use patterns in one platform may reinforce identifiability in others, using data from a holistic, day-in-the-life capture of users of enterprise infrastructure common in offices and academic universities. Practically, the taxonomy and quantification tools developed here could support system audits or compliance reviews, helping identify privacy risks in network traffic previously considered benign.

As encryption increasingly protects content and metadata in various real-world contexts, we expect observable behavioral signals, and network traffic in general, to become a dominant privacy concern as a relatively underexplored gap. By formalizing how these binary observables contribute to anonymity loss, our work provides a foundation for rigorous analysis and defense of privacy in exclusive-use systems.

Ethical Considerations. This research does not involve human subjects or private, real-world user data. To minimize ethical risk, we avoided scraping or collecting third-party data. All evaluations were performed using simulated user activity traces, derived from controlled packet captures in an isolated lab environment to study network traffic sequences. We recognize that the analytical techniques described in this paper could be misused for surveillance or behavioral tracking. To mitigate dual-use concerns, we emphasize transparency, simulation reproducibility, and a focus on defensive applications. Our intent is to support system designers and privacy researchers with tools to understand and mitigate privacy risks, not to facilitate adversarial attacks.

9 Related Work

Behavioral Inference from Metadata. Side-channel attacks have evolved from physical hardware leaks [27, 34, 43] to logical inferences drawn from encrypted or metadata-only traffic. Acar et al.’s Peek-a-Boo attack [1] demonstrated that smart home activities can be inferred through traffic burst patterns, even when payloads are encrypted. Similarly, *IDBleed* [22] showed how binary success/failure signals can expose device affiliations. Hayes and Danezis [37] applied machine learning to identify websites from encrypted traffic traces, while Panchenko et al. [47] used packet attributes measures to assess fingerprinting leakage. These works focus on specific application domains or classification tasks, not anonymity degradation from recurring binary signals.

Metadata Exposure in Messaging Systems. Encrypted messaging platforms often reveal fine-grained metadata such as presence, typing indicators, or delivery receipts. Systems like *Vuvuzela* [60] and *Stadium* [59] introduce mix networks and differential privacy

mechanisms to obscure metadata, routing messages through privacy-preserving infrastructure to achieve strong privacy guarantees. However, these efforts primarily target one-shot exposure or protocol-specific risks, and generally do not address the impact of repeated observables in exclusive-use scenarios.

Quantifying Privacy Loss. QIF offers a formalism for measuring adversary knowledge gain. Smith [55] and Alvim et al. [2] define entropy-based measures of secrecy degradation. Extensions of QIF have been used to evaluate probabilistic programs and quantify leakage under various attacker models [6, 13]. Indistinguishability games [41] provide an alternative framework for reasoning about adversarial advantage in structured guessing scenarios. Our work builds on both perspectives to evaluate how even deterministic binary signals, such as typing or access events, can degrade anonymity over time.

In the context of anonymity systems, Díaz et al. [17] analyze anonymity systems using Shannon entropy to measure uncertainty over sender identities, focusing on mix-based routing models and theoretical anonymity sets. Troncoso and Danezis [58] extend this approach by applying Bayesian inference and Markov Chain Monte Carlo techniques to analyze anonymity degradation in mix networks under probabilistic routing models. Guirat et al. [35] further refine these analyses by studying anonymity under partial network visibility, modeling how incomplete observations affect posterior distributions.

In contrast to routing-centric anonymity analyses, our work applies QIF metrics to exclusive-use systems, where observable binary application-level events are inherently identity-linked and accumulate longitudinally. Rather than modeling path uncertainty in mix networks, we derive posterior distributions directly from behavioral traces and quantify anonymity degradation using both entropy loss and Bayes vulnerability. While entropy captures overall uncertainty reduction, Bayes vulnerability provides an operational measure of attacker success—the probability of a correct one-shot guess—enabling direct comparison of attacker effectiveness and quantitative evaluation of mitigation strategies.

Privacy Attack Taxonomies. Prior work has proposed domain specific taxonomies that organize privacy threats within particular technical contexts. Vidanage et al. survey attacks on privacy preserving record linkage, structuring adversarial goals and linkage vulnerabilities within a protocol specific model [61]. Work on *PRASH* by Bugeja et al. combines an entity-based taxonomy with attack trees and risk scoring tailored to IoT and smart homes environments [11]. Moving toward machine learning privacy attacks, Rigaki et al.’s taxonomy distinguishes inference, extraction, and property leakage threats unique to ML pipelines [53]. Ratnayake et al. [51] review federated learning taxonomies, categorizing threats by data distribution, aggregation, and model update exposure, while Boussada et al. [10] survey contextual privacy attacks in networked systems and classify anonymity, pseudonymity, unlinkability, and unobservability violations across communication protocols. In contrast to these domain constrained frameworks, our taxonomy is defined for exclusive-use systems and is grounded in the attacker’s view of binary observable signals.

Cross-System Fingerprinting. Cross-domain behavioral correlation magnifies the impact of observable signals. Arapinis et al. [5]

Work	Binary Obs.	Taxonomy	Formal Model	Quantification	Cumulative	Entropy/Game	Exclusive-Use	Cross-System
Peek-a-Boo [1]	✓	✗	✗	✓	✗	✗	✓	✗
IDBleed [22]	✓	✗	✓	✗	✗	✗	✓	✗
k-fingerprinting [37]	✗	✗	✗	✓	✗	✗	✗	✗
Alvim [2]	✗	✗	✓	✓	✗	✓	✗	✗
Panchenko [47]	✗	✗	✗	✓	✗	✗	✗	✗
Smith [55]	✗	✗	✓	✓	✗	✓	✗	✗
This Work	✓	✓	✓	✓	✓	✓	✓	✓

Table 3: Comparison of our work with related efforts.

show that protocol-level identifier reuse in mobile systems can reveal session-level behaviors. Reardon et al. [52] catalog how mobile apps expose user behavior through unauthorized data flows, while Wang et al. [62] and Li et al. [44] demonstrate that app usage patterns can be linked to unique user behavior. These works illustrate user fingerprinting feasibility from signals, but do not define a general framework for evaluating binary observables across exclusive-use systems.

Systematization and Scope. Systematization efforts have contextualized these risks across broader domains. Zheng et al. [66] synthesize users’ smart home privacy perceptions, while Reardon et al. [52] document privacy risks in mobile ecosystems. Despite growing awareness of inference risks, prior work lacks a unifying formalism for modeling and quantifying anonymity degradation from low-level signals.

Comparison Summary. Prior work has exposed metadata leakage and behavioral inference in targeted domains, but few efforts address anonymity degradation in exclusive-use systems. Our work formalizes binary interaction signals as a distinct class of leakage, develops an attack taxonomy, and introduces an evaluation framework grounded in entropy and distinguishability. We provide a summary of selected works in Table 3.

10 Conclusion

This paper examined privacy risks in exclusive-use systems, showing that seemingly low-value signals such as presence, typing, and message sends can be digitized into semantic multidimensional vectors that support probabilistic re-identification and a meaningful privacy threat. Our formal model combines anonymity-set reduction, entropy-based analysis, Bayes vulnerability, and indistinguishability games, and our case studies across messaging applications and protocols demonstrate that these binary observables consistently leak behavioral structure. We introduced a taxonomy that organizes exclusive-use privacy attacks by capability, observation method, and leaked information for practical modeling. Our framework provides the first cross-domain foundation for analyzing exclusive-use leakage, enabling defenders to formally model exclusive-use systems under realistic attacker assumptions and to quantitatively evaluate the privacy impact of proposed mitigations using entropy and vulnerability-based metrics.

Acknowledgments

We would like to thank Felix Engelmann for early helpful discussions and we are further grateful to the reviewers for their constructive feedback, which substantially improved the paper. This work was supported in part by the National Security Agency under contract H98230-23-C-0282 and by the National Science Foundation (CNS-2112471 and CNS-2207202).

Generative AI tools were utilized to aid in the revision process, which included improving textual flow and correcting errors in grammar and spelling.

References

- [1] Abbas Acar, Hossein Fereidooni, Tigist Abera, Amit Kumar Sikder, Markus Miettinen, Hidayet Aksu, Mauro Conti, Ahmad-Reza Sadeghi, and Selcuk Uluagac. 2020. Peek-a-boo: i see your smart home activities, even encrypted!. In *Proceedings of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '20)*. Association for Computing Machinery, 207–218. <https://doi.org/10.1145/3395351.3399421>
- [2] M'rio S. Alvim, Kostas Chatzikokolakis, Catuscia Palamidessi, and Geoffrey Smith. 2012. Measuring Information Leakage Using Generalized Gain Functions. In *2012 IEEE 25th Computer Security Foundations Symposium*. 265–279. <https://doi.org/10.1109/CSF.2012.26>
- [3] Sebastian Angel and Srinath Setty. 2016. Unobservable communication over fully untrusted infrastructure. In *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation (Savannah, GA, USA) (OSDI'16)*. USENIX Association, USA, 551–569.
- [4] Apple. 2025. Privacy features when connecting to wireless networks - Apple Support. <https://support.apple.com/guide/security/privacy-features-connecting-wireless-networks-secb9cb3140c/web>. Accessed 2025-11-22.
- [5] Myrto Arapinis, Loretta Mancini, Eike Ritter, Mark Ryan, Nico Golde, Kevin Redon, and Ravishankar Borgaonkar. 2012. New privacy issues in mobile telephony: fix and verification. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security (Raleigh, North Carolina, USA) (CCS '12)*. Association for Computing Machinery, New York, NY, USA, 205–216. <https://doi.org/10.1145/2382196.2382221>
- [6] Michael Backes, Boris Köpf, and Andrey Rybalchenko. 2009. Automatic Discovery and Quantification of Information Leaks. In *2009 30th IEEE Symposium on Security and Privacy*. 141–153. <https://doi.org/10.1109/SP.2009.18>
- [7] Ludovic Barman, Moshe Kol, David Lazar, Yossi Gilad, and Nickolai Zeldovich. 2022. Groove: Flexible Metadata-Private Messaging. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*. USENIX Association, Carlsbad, CA, 735–750. <https://www.usenix.org/conference/osdi22/presentation/barman>
- [8] Nikita Borisov, George Danezis, and Ian Goldberg. 2015. DP5: A Private Presence Service. *Proceedings on Privacy Enhancing Technologies* 2015 (06 2015). <https://doi.org/10.1515/popets-2015-0008>
- [9] Ioana Boureanu, Tom Chothia, Alexandre Debant, and Stéphanie Delaune. 2020. Security Analysis and Implementation of Relay-Resistant Contactless Payments. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security (Virtual Event, USA) (CCS '20)*. Association for Computing Machinery, New York, NY, USA, 879–898. <https://doi.org/10.1145/3372297.3417235>
- [10] Rihab Boussada, Mohamed Elhoucine Elhdhili, and Leila Azouz Saidane. 2016. A survey on privacy: Terminology, mechanisms and attacks. In *2016 IEEE/ACS 13th International Conference of Computer Systems and Applications (AICCSA)*. 1–7. <https://doi.org/10.1109/AICCSA.2016.7945804>
- [11] Joseph Bugeja, Andreas Jacobsson, and Paul Davidsson. 2021. PRASH: A Framework for Privacy Risk Analysis of Smart Homes. *Sensors* 21, 19 (2021). <https://doi.org/10.3390/s21196399>
- [12] Xiang Cai, Xin Cheng Zhang, Brijesh Joshi, and Rob Johnson. 2012. Touching from a distance: website fingerprinting attacks and defenses. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security (Raleigh, North Carolina, USA) (CCS '12)*. Association for Computing Machinery, New York, NY, USA, 605–616. <https://doi.org/10.1145/2382196.2382260>
- [13] Konstantinos Chatzikokolakis, Tom Chothia, and Apratim Guha. 2010. Statistical measurement of information leakage. In *Proceedings of the 16th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (Paphos, Cyprus) (TACAS '10)*. Springer-Verlag, Berlin, Heidelberg, 390–404. https://doi.org/10.1007/978-3-642-12002-2_33
- [14] Michael Clarkson, Andrew Myers, and Fred Schneider. 2009. Quantifying information flow with beliefs. *Journal of Computer Security* 17 (10 2009), 655–701. <https://doi.org/10.3233/JCS-2009-0353>
- [15] Henry Corrigan-Gibbs, Dan Boneh, and David Mazières. 2015. Riposte: An Anonymous Messaging System Handling Millions of Users. In *2015 IEEE Symposium on Security and Privacy*. 321–338. <https://doi.org/10.1109/SP.2015.27>
- [16] Trisha Datta, Noah Aporthe, and Nick Feamster. 2018. A Developer-Friendly Library for Smart Home IoT Privacy-Preserving Traffic Obfuscation. In *Proceedings of the 2018 Workshop on IoT Security and Privacy (Budapest, Hungary) (IoT S&P '18)*. Association for Computing Machinery, New York, NY, USA, 43–48. <https://doi.org/10.1145/3229565.3229567>
- [17] Claudia Diaz, Stefaan Seys, Joris Claessens, and Bart Preneel. 2003. Towards Measuring Anonymity. In *Privacy Enhancing Technologies*, Roger Dingledine and Paul Syverson (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 54–68.
- [18] Whitfield Diffie and Martin E Hellman. 1976. New directions in cryptography. *IEEE Transactions on Information Theory* 22, 6 (1976), 644–654. <https://doi.org/10.1109/TIT.1976.1055638>
- [19] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 265–284.
- [20] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Found. Trends Theor. Comput. Sci.* 9, 3–4 (Aug. 2014), 211–407. <https://doi.org/10.1561/04000000042>
- [21] Christopher Ellis, Hao Huang Wen, Zhiqiang Lin, and Anish Arora. 2022. Replay (Far) Away: Exploiting and Fixing Google/Apple Exposure Notification Contact Tracing. *Proceedings on Privacy Enhancing Technologies* 2022 (10 2022), 727–745. <https://doi.org/10.56553/popets-2022-0130>
- [22] Christopher Ellis, Yue Zhang, Mohit Kumar Jangid, Shixuan Zhao, and Zhiqiang Lin. 2025. Deanonimizing Device Identities via Side-channel Attacks in Exclusive-use IoTs & Mitigation. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*. Internet Society.
- [23] J. Elson and D. Estrin. 2001. Random, Ephemeral Transaction Identifiers in dynamic sensor networks. In *Proceedings 21st International Conference on Distributed Computing Systems*. 459–468. <https://doi.org/10.1109/ICDSC.2001.918976>
- [24] Xabier Etxezarreta, I naki Garitano, Mikel Iturbe, and Urko Zurutuza. 2024. Low delay network attributes randomization to proactively mitigate reconnaissance attacks in industrial control systems. *Wireless Networks* 30, 6 (2024), 5077–5091. <https://doi.org/10.1007/s11276-022-03212-5>
- [25] European Union. 2018. General Data Protection Regulation. (2018). <https://gdpr.eu/>.
- [26] Stephan A. Fahrenkrog-Petersen, Han van der Aa, and Matthias Weidlich. 2023. Optimal event log sanitization for privacy-preserving process mining. *Data & Knowledge Engineering* 145 (2023), 102175. <https://doi.org/10.1016/j.datak.2023.102175>
- [27] Karine Gandolfi, Christophe Mourtel, and Francis Olivier. 2001. Electromagnetic Analysis: Concrete Results. In *Cryptographic Hardware and Embedded Systems — CHES 2001*, Çetin K. Koç, David Naccache, and Christof Paar (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 251–261.
- [28] Gabriel Karl Gegenhuber, Philipp Frenzel, Maximilian Günther, Johanna Ullrich, and Aljosha Judmayer. 2025. Hey there! You are using WhatsApp: Enumerating Three Billion Accounts for Security and Privacy. <https://doi.org/10.48550/arXiv.2511.20252>
- [29] Patric Genfer and Uwe Zdun. 2022. Avoiding Excessive Data Exposure Through Microservice APIs. In *Software Architecture: 16th European Conference, ECSA 2022, Prague, Czech Republic, September 19–23, 2022, Proceedings (Prague, Czech Republic)*. Springer-Verlag, Berlin, Heidelberg, 3–18. https://doi.org/10.1007/978-3-031-16697-6_1
- [30] Oded Goldreich. 2004. *Foundations of Cryptography: Volume 2, Basic Applications*. Cambridge University Press.
- [31] Shafi Goldwasser and Silvio Micali. 1984. Probabilistic encryption. *J. Comput. System Sci.* 28, 2 (1984), 270–299. [https://doi.org/10.1016/0022-0000\(84\)90070-9](https://doi.org/10.1016/0022-0000(84)90070-9)
- [32] Google. 2025. MAC randomization behavior - Android Open Source Project. <https://source.android.com/docs/core/connect/wifi-mac-randomization-behavior>. Accessed 2025-11-22.
- [33] Tim Grube, Markus Thummerer, Jörg Daubert, and Max Mühlhäuser. 2017. Cover Traffic: A Trade of Anonymity and Efficiency. In *Security and Trust Management*, Giovanni Livraga and Chris Mitchell (Eds.). Springer International Publishing, Cham, 213–223.
- [34] Daniel Gruss, Clémentine Maurice, Klaus Wagner, and Stefan Mangard. 2016. Flush+Flush: A Fast and Stealthy Cache Attack. In *Proceedings of the 13th International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment - Volume 9721 (San Sebastián, Spain) (DIMVA 2016)*. Springer-Verlag, Berlin, Heidelberg, 279–299. https://doi.org/10.1007/978-3-319-40667-1_14
- [35] Iness Ben Guirat, Claudia Diaz, Karim Eldefrawy, and Hadas Zeilberger. 2024. Traffic Analysis by Adversaries with Partial Visibility. In *Computer Security – ESORICS 2023*, Gene Tsudik, Mauro Conti, Kaitai Liang, and Georgios Smaragdakis (Eds.). Springer Nature Switzerland, Cham, 338–358.
- [36] Avinatan Hassidim, Yossi Matias, Moti Yung, and Alon Ziv. 2016. Ephemeral Identifiers: Mitigating Tracking & Spoofing Threats to BLE Beacons. <https://api.semanticscholar.org/CorpusID:26483080>
- [37] Jamie Hayes and George Danezis. 2016. k-fingerprinting: A Robust Scalable Website Fingerprinting Technique. In *25th USENIX Security Symposium (USENIX Security and Privacy)*. 321–338. <https://doi.org/10.1109/SP.2015.27>

- Security 16). USENIX Association, Austin, TX, 1187–1203. <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/hayes>
- [38] Jason I. Hong and James A. Landay. 2004. An architecture for privacy-sensitive ubiquitous computing. In *Proceedings of the 2nd International Conference on Mobile Systems, Applications, and Services* (Boston, MA, USA) (*MobiSys '04*). Association for Computing Machinery, New York, NY, USA, 177–189. <https://doi.org/10.1145/990064.990087>
- [39] Xintao Huan, Kaitao Miao, Wen Chen, Pengyi Jia, and Han Hu. 2024. Kerra: An Internet of Things Wireless Key Generation Resistant to Replay Attacks. *IEEE Internet of Things Journal* 11, 17 (2024), 29035–29048. <https://doi.org/10.1109/JIOT.2024.3406702>
- [40] Pierre-Marie Junges, Jérôme François, and Olivier Festor. 2019. Passive Inference of User Actions through IoT Gateway Encrypted Traffic Analysis. In *2019 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*. 7–12.
- [41] Jonathan Katz and Yehuda Lindell. 2014. *Introduction to Modern Cryptography* (2 ed.). Chapman and Hall/CRC. <https://doi.org/10.1201/b17668>
- [42] Paul Kocher, Joshua Jaffe, and Benjamin Jun. 1999. Differential Power Analysis. In *Advances in Cryptology — CRYPTO '99*, Michael Wiener (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 388–397.
- [43] Paul C. Kocher. 1996. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In *Advances in Cryptology — CRYPTO '96*, Neal Koblitz (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 104–113.
- [44] Tong Li, Yong Li, Mingyang Zhang, Sasu Tarkoma, and Pan Hui. 2023. You Are How You Use Apps: User Profiling Based on Spatiotemporal App Usage Behavior. *ACM Trans. Intell. Syst. Technol.* 14, 4, Article 71 (July 2023), 21 pages. <https://doi.org/10.1145/3597212>
- [45] Xipei Luo, Jing Wang, Qiwei Shen, Jingyu Wang, and Qi Qi. 2017. User behavior analysis based on user interest by web log mining. In *2017 27th International Telecommunication Networks and Applications Conference (ITNAC)*. 1–5. <https://doi.org/10.1109/ATNAC.2017.8215435>
- [46] MITRE. 2025. MITRE ATT&CK. <https://attack.mitre.org/>. Accessed 2025-11-21.
- [47] Andriy Panchenko, Fabian Lanze, Andreas Zinnen, Martin Henze, Jan Pennekamp, Klaus Wehrle, and Thomas Engel. 2016. Website Fingerprinting at Internet Scale. <https://doi.org/10.14722/ndss.2016.23477>
- [48] Andreas Pfitzmann and Marit Hansen. 2010. A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management. *Technical report, TU Dresden* (2010). https://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf
- [49] Phillip Porras and Vitaly Shmatikov. 2006. Large-scale collection and sanitization of network security data: risks and challenges. In *Proceedings of the 2006 Workshop on New Security Paradigms* (Germany) (*NSPW '06*). Association for Computing Machinery, New York, NY, USA, 57–64. <https://doi.org/10.1145/1278940.1278949>
- [50] Apostolos Pyrgelis, Carmela Troncoso, and Emiliano De Cristofaro. 2017. Knock Knock, Who's There? Membership Inference on Aggregate Location Data. [arXiv:1708.06145 \[cs.CR\]](https://arxiv.org/abs/1708.06145) <https://arxiv.org/abs/1708.06145>
- [51] Hashan Ratnayake, Lin Chen, and Xiaofeng Ding. 2023. A review of federated learning: taxonomy, privacy and future directions. *J. Intell. Inf. Syst.* 61, 3 (July 2023), 923–949. <https://doi.org/10.1007/s10844-023-00797-x>
- [52] Joel Reardon, Álvaro Feal, Primal Wijesekera, Amit Elazari Bar On, Narseo Vallina-Rodriguez, and Serge Egelman. 2019. 50 Ways to Leak Your Data: An Exploration of Apps' Circumvention of the Android Permissions System. In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, 603–620.
- [53] Maria Rigaki and Sebastian Garcia. 2023. A Survey of Privacy Attacks in Machine Learning. *ACM Comput. Surv.* 56, 4, Article 101 (Nov. 2023), 34 pages. <https://doi.org/10.1145/3624010>
- [54] Bluetooth SIG. 2025. Bluetooth Core Specification. <https://www.bluetooth.com/wp-content/uploads/Files/Specification/HTML/Core-60/out/en/index-en.html>. Accessed 2025-05-29.
- [55] Geoffrey Smith. 2009. On the Foundations of Quantitative Information Flow. In *Foundations of Software Science and Computational Structures*, Luca de Alfaro (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 288–302.
- [56] Ye Tao, Shuaitong Guo, Cao Shi, and Dianhui Chu. 2020. User Behavior Analysis by Cross-Domain Log Data Fusion. *IEEE Access* 8 (2020), 400–406. <https://doi.org/10.1109/ACCESS.2019.2961769>
- [57] Time Is Ltd. 2021. Transparency is Key to Making Slack Work Within Your Organization. <https://www.timeisltd.com/post/transparency-is-key-to-making-slack-work-within-your-organization>. Accessed May 2025.
- [58] Carmela Troncoso and George Danezis. 2009. The bayesian traffic analysis of mix networks. In *Proceedings of the 16th ACM Conference on Computer and Communications Security* (Chicago, Illinois, USA) (*CCS '09*). Association for Computing Machinery, New York, NY, USA, 369–379. <https://doi.org/10.1145/1653662.1653707>
- [59] Nirvan Tyagi, Yossi Gilad, Derek Leung, Matei Zaharia, and Nickolai Zeldovich. 2017. Stadium: A Distributed Metadata-Private Messaging System. In *Proceedings of the 26th Symposium on Operating Systems Principles* (Shanghai, China) (*SOSP '17*). Association for Computing Machinery, New York, NY, USA, 423–440. <https://doi.org/10.1145/3132747.3132783>
- [60] Jelle van den Hooff, David Lazar, Matei Zaharia, and Nickolai Zeldovich. 2015. Vuvuzela: scalable private messaging resistant to traffic analysis. In *Proceedings of the 25th Symposium on Operating Systems Principles* (Monterey, California) (*SOSP '15*). Association for Computing Machinery, New York, NY, USA, 137–152. <https://doi.org/10.1145/2815400.2815417>
- [61] Anushka Vidanage, Thilina Ranbaduge, Peter Christen, and Rainer Schnell. 2022. A Taxonomy of Attacks on Privacy-Preserving Record Linkage. *Journal of Privacy and Confidentiality* 12, 1 (Jul. 2022). <https://doi.org/10.29012/jpc.764>
- [62] Qinglong Wang, Amir Yahyavi, Bettina Kemme, and Wenbo He. 2015. I know what you did on your smartphone: Inferring app usage over encrypted data traffic. In *2015 IEEE Conference on Communications and Network Security (CNS)*. 433–441. <https://doi.org/10.1109/CNS.2015.7346855>
- [63] Tao Wang, Xiang Cai, Rishab Nithyanand, Rob Johnson, and Ian Goldberg. 2014. Effective attacks and provable defenses for website fingerprinting. In *Proceedings of the 23rd USENIX Conference on Security Symposium* (San Diego, CA) (*SEC'14*). USENIX Association, USA, 143–157.
- [64] Sijie Xiong, Anand D. Sarwate, and Narayan B. Mandayam. 2022. Network Traffic Shaping for Enhancing Privacy in IoT Systems. *IEEE/ACM Transactions on Networking* 30, 3 (2022), 1162–1177. <https://doi.org/10.1109/TNET.2021.3140174>
- [65] Yue Zhang and Zhiqiang Lin. 2022. When Good Becomes Evil: Tracking Bluetooth Low Energy Devices via Allowlist-based Side Channel and Its Countermeasure. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security* (Los Angeles, CA, USA) (*CCS '22*). Association for Computing Machinery, New York, NY, USA, 3181–3194. <https://doi.org/10.1145/3548606.3559372>
- [66] Serena Zheng, Noah Aphorpe, Marshini Chetty, and Nick Feamster. 2018. User Perceptions of Smart Home IoT Privacy. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 200 (Nov. 2018), 20 pages. <https://doi.org/10.1145/3274469>

A Additional Figures

This appendix includes additional evaluation results and visualizations.




Figure 9: Raster timeline of simulated user activity. Rows represent users, and markers indicate presence, typing, and message events over time.