Classes of Cyber Physical System Observation Privacy Techniques

Chathuranga Sampath Kalutharage
School of Computing and Communications
Lancaster University
Lancaster, United Kingdom
c.kalutharage@lancaster.ac.uk

Matthew Bradbury
School of Computing and Communications
Lancaster University
Lancaster, United Kingdom
m.s.bradbury@lancaster.ac.uk

Abstract

Content privacy protects data confidentiality through encryption during storage and transmission or via privacy-preserving transformations. However, as Cyber Physical Systems act in a physical space they are also vulnerable to direct observations on the actions taken by the system, its state, and the context in which those actions occur. In such a case, content privacy is insufficient (but often necessary) to provide privacy and additional techniques are required. Across multiple areas, many approaches have been taken to reduce information loss to direct adversary observations, however, research across these different cyber physical system domains has not typically interacted. Therefore, this work deconstructs and systematises existing context privacy techniques into three classes: (i) Add Noise, (ii) Decorrelate, and (iii) Change Observability. We also speculate on class of (iv) Make Sensitive Commonplace techniques. Each class captures a distinct strategy to mitigate information leakage from adversary observations. We illustrate this taxonomy using an example where an adversary observes public transport interactions, showing how each class can be instantiated with a representative context privacy technique. We classify a broad range of past work protecting cyber physical systems from observing adversaries and identify potential gaps in areas where classes of techniques have not been explored in depth.

CCS Concepts

• **Security and privacy**; • **Computer systems organization** → *Embedded and cyber-physical systems*;

Keywords

Observation Privacy, Technique Classes

ACM Reference Format:
Chathuranga Sampath Kalutharage and Matthew Bradbury. 2026. Classes of Cyber Physical System Observation Privacy Techniques. In *12th ACM Cyber-Physical System Security Workshop (CPSS '26), June 1–5, 2026, Bangalore, India.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3775042.3807879

1 Introduction

Protecting the privacy of data content is a vital aspect of a system's security. It is necessary to ensure that only authenticated users are able to access data at rest and that confidentiality is provided for data in motion. Many techniques exist to provide this confidentiality of data; however, there is much that can be learned from systems by observing the *actions* that systems perform and the *context* in which they occur. This is especially important for Cyber Physical Systems (CPSs) as they exist in a physical space which provides additional vectors for adversaries to observe those systems. For example, this might include the tracking of connected vehicles [8], observing routing patterns to locate valuable assets in Wireless Sensor Networks [11], or inferring information about users from their energy consumption [23].

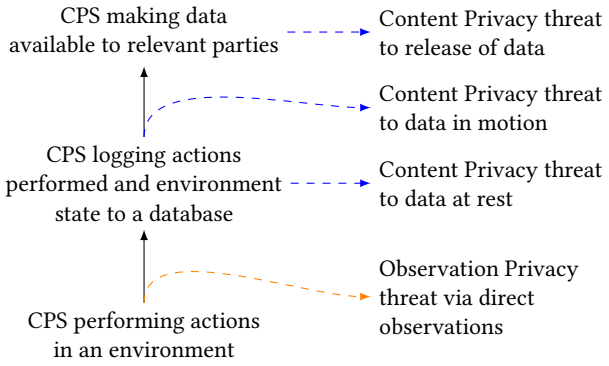

Figure 1: Stages of privacy threats

In each of these example cyber physical systems, data confidentiality techniques such as encryption are essential, however, the privacy threat of observing information about the system remains because the observations on the system occur before this information becomes data (as shown in Figure 1) which can be protected with typical privacy techniques. For connected vehicles, the tracking of identifiers as the vehicle changes location leads to inference about where the driver lives, works and facilitates predictions of this information in the future [32]. For Wireless Sensor Networks, it is an adversary observing the direction from which messages are sent that leads to privacy loss [27]. For energy consumption, it is the need to accurately report energy usage for billing that allows pattern of life inference [23]. Each of these privacy problems has an adversary take a different approach to observe the system and the typical confidentiality approach of encryption is insufficient to mitigate to the threat. As is often the case with such side-channel attacks (e.g., power side-channel attacks on reconstructing 3D printed models [18]). This means that different approaches need to be taken where system behaviour is changed to mitigate the privacy threat.

A problem with past work is that both the privacy threat and the technique developed to address that privacy threat are developed within the specific CPS domain from which the problem originates.

This work is licensed under a Creative Commons Attribution 4.0 International License.
CPSS '26, Bangalore, India
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2313-1/26/06
https://doi.org/10.1145/3775042.3807879

There is a lack of transference of both awareness of issues and the techniques from one CPS domain to another. To address this issue and to identify scope for the development of new types of observation privacy techniques in different CPS domains, in this paper we categorise existing techniques that mitigate these privacy threats into three categories: (i) Add Noise, (ii) Decorrelate, and (iii) Change Observability, and introduce a fourth speculative class (iv) Make Sensitive Commonplace later in the paper. Our focus is on system-level strategies which require a change in the system (e.g., in terms of its behaviour) that reduce an adversary's ability to link observations or infer sensitive behaviour. By establishing this categorisation, we aim to enable comparison across different approaches takes by different CPS and identify where types of techniques could be applied to CPSs that have not previously used them to address observation privacy threats. The contributions of this work are as follows:

(1) We propose a taxonomy that organises observation privacy techniques into three classes and include a fourth speculative class in the Discussion (Section 6).
(2) We present a representative example system of a public transportation network which is used to demonstrate how each class of technique could be applied in a common setting.
(3) We demonstrate the relevance of these techniques beyond this example by identifying their applicability to different cyber physical systems and other areas.

The rest of this paper is structured as follows: Firstly, in Section 2 the common threat model used across the different technique classes is presented. Next, in Section 3 we introduce three classes of context privacy techniques and explain our assumptions about the threat actor that the technique class is attempting to mitigate against. We then present our public transport scenario in Section 4, where observations are made on interactions between users and the ticket barriers cyber physical system. Using our four classes from Section 3, we explore representatives techniques to demonstrate how simple mitigations fall into these classes. This is followed by a review of relevant literature in Section 5 where we identify which classes these techniques fall under to identify what types of approaches have been used in different areas. In Section 6 we discussion of these technique classes and present a fourth speculative class. The paper concludes with reflections on the broader implications of this work and avenues for further investigation in Section 7.

2 Threat Model

In this section we summarise a general threat model for data privacy, which is used to situate the difference between these commonly considered data privacy threats and observational threats to cyber physical systems.

2.1 Data Privacy

Data Privacy is typically considered in four use-cases. Firstly, protecting confidentiality of data in transit. There is typically an adversary either intercepting or eavesdropping communications. The adversary may be remote (e.g., to intercept internet traffic) or local (e.g., eavesdropping wireless communication). Secondly, protecting confidentiality of data at rest where an adversary may attempt to gain higher privileges in order to access to access data whose access is restricted. Alternatively, adversaries may undertake physical attacks in order to steal physical storage media. Thirdly, undertaking a computation over private data across multiple parties without revealing the data to the other parties. Finally, providing bounds on privacy loss when releasing of data. This threat is different, as there is a desire to release data, however, the inferences that can be made by an adversary on the released data needs to be constrained.

2.2 Observation Privacy

In this work, we consider an adversary that is capable of observing a cyber physical system, this includes observing the state of the system, the actions taken, and the context of these observations.

Observable actions could include movement, actuation, or communication patterns. Observable states could include location, physical position, physical condition, and other characteristics of the system. These observations may be made by visual, acoustic, wireless, monitoring resource consumption, or various other modes. If these observations are made on communications, then it is assumed that these communications will be encrypted to protect data in transit and privacy loss will occur via contextual information (e.g., timing, packet size, signal strength, transmission direction, etc.) or packet header information that is not suitable to be encrypted for the network to function.

The adversary is assumed to be passive and observes the system over time. An active adversary that interacts with the system to trigger states and actions to be observed has the potential to be detected by the system being observed and thus existing work on intrusion detection and prevention techniques for cyber physical systems can be used [28]. By analysing patterns across multiple observations, the adversary can infer sensitive information such as user behaviour, routines, or location [26].

In this setting, content privacy techniques such as encryption are necessary but insufficient to mitigate these observation threats. Without content privacy techniques more information will be revealed to an observing adversary. However, because information leakage occurs through observable states and actions, which include physical states that cannot be encrypted, additional techniques are required to reduce what an adversary can learn from observations.

3 Observation Privacy Classes

This section introduces three classes of observation privacy techniques. These techniques are motivated by the need to protect information about cyber physical systems such as why an action occurred, when an action occurred, or under what circumstances an action is taken. These classes are specifically designed to address a powerful threat model in which an adversary is capable of directly observing system interactions and inferring sensitive contextual details. These classes group techniques by the approach taken to reduce information loss to the observing adversary. While each of these three classes assume the common threat model defined in Section 2.2 they have specific refinements of this threat model which techniques in this class mitigate.

3.1 Add Noise

Noise-based techniques introduce randomness to observable actions, making it more difficult for an adversary to distinguish meaningful patterns. This could involve introducing new unnecessary activities to the system, or changing how activities are performed. Here the intuition is that by introducing behavioural noise, the adversary's certainty about the system, its state and intentions is decreased.

This class of observation privacy techniques contrasts with data privacy techniques where noise may be added to datasets (e.g., via differential privacy [15]) to provide bounds on privacy loss. However, for this observational privacy problem, there is no potential to add noise to a dataset as the adversary is directly observing the system and generating the observational data itself which cannot be transformed, so only the behaviour of the system can be adjusted.

The adversary may be trying to infer patterns in the behaviour of the system over time in order to predict future behaviour, or to identify some sensitive aspect of the system. For example, this might mean cyber physical systems need to: take additional, unnecessary journeys [39], send dummy data packets in a network [7], or generate other fake events [40]. The added noise increases the difficulty of statistical inference, making it computationally expensive or practically infeasible [17] for an observer to extract accurate insights about the cyber physical system.

3.2 Decorrelate

Decorrelation techniques break the linkage between observations made by an adversary and an underlying sensitive context. The goal is to limit or prevent an adversary from connecting events over time to infer motivation or identity.

The adversary is an observer who is able to observe time, location, identifiers, or other actions and link them to construct activity patterns. Techniques in this category to mitigate this threat may anonymize persistent identifiers [8] to mitigate long-term tracking of the system, randomise communication timings, or use additional equipment to hide activities [34]. By eliminating or weakening the correlation between actions, the system maintains a disconnect between its motivations or intent and the actions it takes.

3.3 Change Observability

Sometimes the most effective technique is to restrict what an adversary can observe in the first place. These techniques that change observability seek to limit or prevent an adversary from being able to make observations on the system.

For example, this might be to change mechanisms via which the system interacts to entirely different mechanisms which the adversary does not have the capability to observe. For example, this could be to move from using wireless communication which is vulnerable to eavesdropping to wired communication. Alternatively, the system to cease interacting with the components via which the adversary makes its observations. In general, this is the most challenging technique for systems to employ without impacting a cyber-physical system's ability to perform their functions.

4 Example Scenario

In this section we present an example scenario involving the use of public transport to highlight the differences in the four classes of techniques. The adversary observes users interacting with the public transport system. This includes when and where users enter and exit the system. By collecting these observations over time, the adversary can infer sensitive information, such as frequently visited locations or unusual behaviour. This scenario is chosen as it was feasible to implement each of the four classes of techniques, which is not always realistic in other areas. These techniques are intended to be simple and illustrative, they are not intended to be state-of-the-art solutions to these privacy threats.

The public transport network is a directed graph $G = (V, E)$ where V represents stops in the network and $E \subseteq V \times V$ is the set of ordered pairs representing a single hop that can be taken in the public transport. A journey J is a sequence of stops $J = J_0 \cdot J_1 \cdots J_n$ where $J_i \in V$. J_0 represents an entry into the public transport network, J_n represents an exit from the public transport network and all other stops J_i represent an exit and re-entry into the public transport network. So here a journey is not the route taken through the public transport network, but the points where users enter/exit from the system.

As this work focuses on an observation privacy problem, we assume that the underlying infrastructure, such as the public transport network cannot be modified by the user to provide privacy. This includes the physical layout of stops and routes, as well as system-level components such as the payment protocol. Specifically, we assume that traditional data privacy techniques (e.g., encryption of transmitted data) are already in place where appropriate.

For this example, we consider privacy leakage that arises from where and when a user enters or exits the system. This means the threat actor performs their observations at the point users need to enter or exit the system. We assume that users will use a smartcard to *tap-in* to start a journey and *tap-out* to end a journey. The smartcard will use a fixed identity allowing users' activity to be linked over time. The adversary could use multiple approaches to perform this observation. The first is if the public transportation system itself was the threat actor the user was attempting to limit the usefulness of their observations to. This is a similar threat model to that used in energy consumption privacy where the smart meter is the point of observation [23, 34, 35, 41]. Alternatively, adversaries may eavesdrop on the wireless communications between the smartcard and the reader. As this will typically be implemented using RFID with a typical range of 10 cm, it means that an eavesdropper will need to be within at most a few meters to eavesdrop communications [20]. Each of the categories of techniques will present a refinement of this threat model, as different techniques will be used to address different observational privacy threats.

4.1 Add Noise

4.1.1 Refined Threat Model. The adversary collects many journeys over time, through smartcard tap in/out. This facilities pattern of life analysis and allows for the prediction of future behaviour and the importance of certain locations (e.g., home, work, etc).

4.1.2 Privacy Goal. The goal with this class of technique is to introduce behavioural noise via unnecessary actions to obscure the

genuine actions taken. This decreases the certainty of the adversary in what the genuine actions were, but it is not to prevent genuine actions from being observed by the adversary.

4.1.3 Example Technique. This illustrative technique which introduces noise to user behaviour aims to obscure which location the user visited by introducing additional stops to the journey which the user does not need to visit. Suppose a user wishes to travel from node $a \in V$ to their true destination $b \in V$. Rather than taking the direct path $a \to b$, the user introduces obfuscation by visiting a sequence of randomly chosen locations. The modified path becomes:

$$a \to c_1 \to c_2 \to \cdots \to c_n \to b \to c_{n+1} \to\!\to c_{n+2} \to \cdots \to c_m$$

Figure 2: Illustration of how a journey could be obfuscated by inserting intermediate nodes before and after the true destination b.

In Figure 2 c_1 to c_n and c_{n+1} to c_m are randomly selected intermediate visits, where the user may optionally wait for a short time at each before continuing the journey.

To evaluate how this strategy reduces the adversary's belief in the true destination, we model the probability of successful inference as inversely proportional to the number of plausible endpoints in the journey. Assuming that each of the $n + 1$ final nodes (including the true destination) is equally likely from the adversary's viewpoint, we define:

$$\Pr = \frac{1}{n+1}.$$

This expression reflects a uniform uncertainty model, where each candidate destination has equal probability in the adversary's estimation. As the number of possible destinations n increase, the inference probability decreases. However, a downside is the increase in cost (both monetary and time) to make these additional visits.

4.2 Decorrelate

4.2.1 Refined Threat Model. The adversary observes smartcard identities at entry and exit points. For complex journeys there will be multiple entry and exit points which can be linked via the identities, allowing the adversary to reconstruct the entire journey taken by a user.

4.2.2 Privacy Goal. In this technique the adversary aims to observe the entire journey, so the privacy goal is to reduce the ability for the adversary to observe this journey. This is done by decorrelating actions of a user entering/exiting the public transport system and the action of them tapping a smartcard to trigger the entrance/exiting.

4.2.3 Example Technique. This technique reduces the correlation between observations that an adversary makes and the underling action taken by the user. Users tap-in and tap-out using smartcards with persistent identifiers. The adversary can observe these identifiers and link entrances and exits over time to obtain the full journey made by the user. To mitigate this, users alter they interaction with the infrastructure by swapping their smartcards with other users throughout their journeys when at the same location. Consider two users u_1 and u_2 who undertake the following journeys:

- u_1 tapped-in at V_1, visited V_2, and tapped-out at V_1, and
- u_2 tapped-in at V_3, visited V_2, and tapped-out at V_3.

Without any mitigation, the adversary would observe the complete routes of both users. When applying this technique where u_1 swaps their smartcard with u_2 at V_2, the adversary now sees the following:

- u_1 tapped-in at V_1, visited V_2, and tapped-out at V_3, and
- u_2 tapped-in at V_3, visited V_2, and tapped-out at V_1.

The key difference is that the observed exit stop in the journey is no longer correct for both users.

To illustrate this technique, we model a 14-node public transit network with $N = 10$ users, a journey length of $T = 10$, and 1000 repeated simulations. Users 0 and 1 begin at node V_1, users 2–4 at V_2, and users 5–9 at randomly assigned starting nodes. At each non-final time step ($t < T - 1$), users transition to a randomly selected node drawn from a uniform distribution over all possible destinations. User 0 does not swap identities to demonstrate the impact of the technique not being in use. Other users can swap identities when they are located at the same stop. An example set of journeys is shown in Figure 3, where users 0 to 9 are assigned identities A to J respectively, and the graph shows how identities cease to be correlated with their initial user activity after swaps.

To illustrate how different these observed behaviours are, we quantify the Jensen-Shannon Divergence (JSD) [38] between two probability distributions obtained from the simulation:

(1) $P_{u,t}$ — which is the probability distribution of where the user u visits at time t, and
(2) $Q_{u,t}$ — which is the probability distribution of where the adversary observes the identity of user u visits at time t.

This is shown in Figure 4, where the users (1 to 9) who engage with swapping identities when they are co-located have an increased divergence between the actions taken and observations made at later time steps compared to the user (0) who does not engage in identity swapping.

4.3 Change Observability

4.3.1 Refined Threat Model. For this class, the refined threat model is the same as Decorrelate where an adversary attempts to reconstruct an entire journey.

4.3.2 Privacy Goal. The adversary is capable of observing all interactions made with the system and the user cannot change how the infrastructure records or timestamps journeys. So if the user wishes to prevent observation of their actions, the only approach is to change how the system is interacted with.

4.3.3 Example Technique. In this example technique, we change the observability of actions by changing how the user interacts with the ticket barriers. Instead of tapping-in and tapping-out a user will tailgate or otherwise bypass the need to provide a smartcard in order to enter or exit the public transport.

To highlight the effect of this technique, we implemented a simulation on a directed transit network comprising 14 nodes. Ten users each undertook a 10-step journey, where the movement at

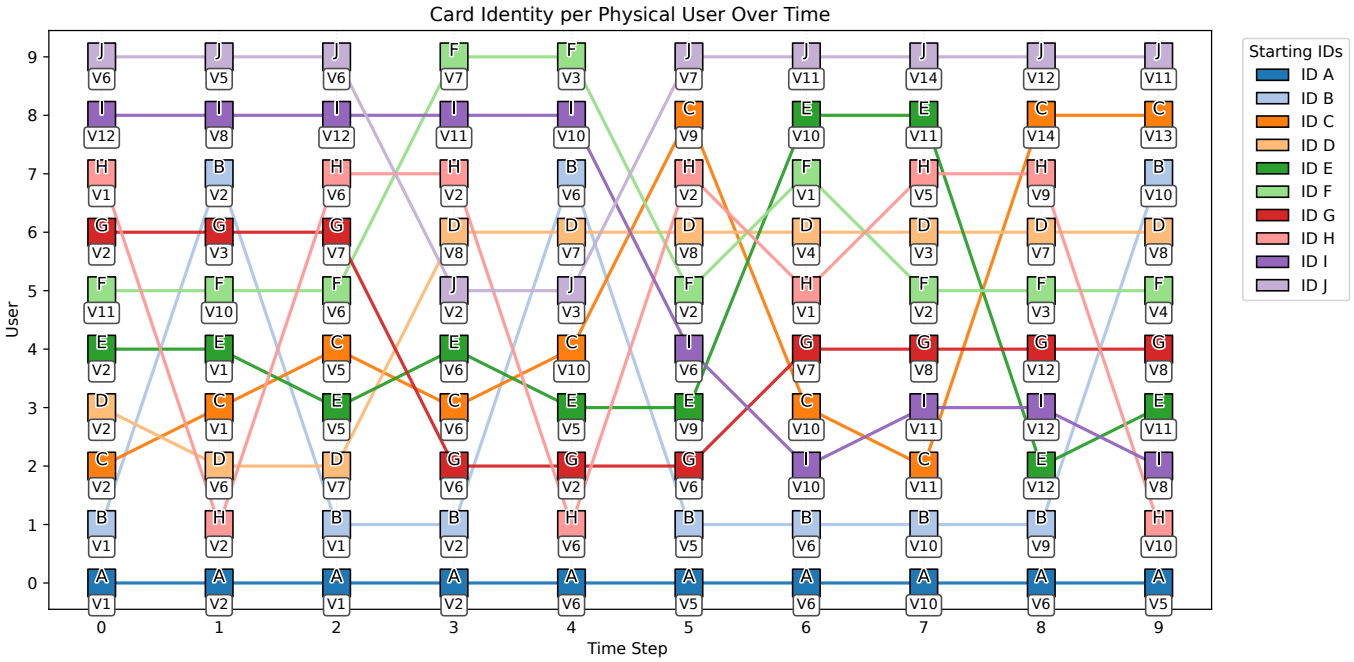

Figure 3: Card swap when users are at the same location.

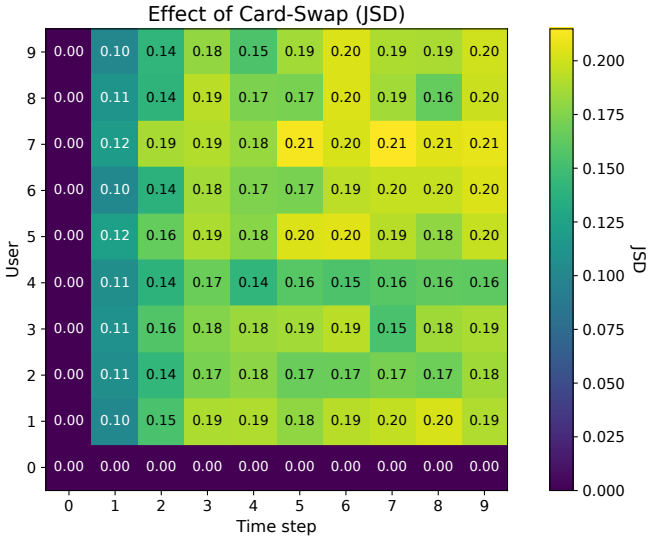

Figure 4: Jensen–Shannon divergence between the actual route taken by a user and the route the adversary observes.

each time step was randomly chosen, and the entire process was repeated 100 times to build a probability distribution of stops over time. To model users circumventing observation by the adversary, we randomly selected three visits for which the stop by the users was unobservable to the adversary (except User 0, whose actions were always observable as a point of comparison). Figure 5 shows a single instance from the 100 repeats where the green line represents the true journey taken and the red and orange lines are adversary observations under two different threat models, either (1) where the adversary is aware that a stop was made (orange) or (2) where the adversary cannot tell that the stop was made (red).

We compute the Jensen-Shannon Divergence (JSD) [38] between the genuine user actions $P_{u,t}$ and observations made by the adversary $Q_{u,t}$ considering the two threat models from before: either the adversary is aware a stop was made (shown in Figure 6a) or the adversary was unaware (shown in Figure 6b). When the adversary is aware of when a stop occurs we see full divergence between the true journey and the adversary observations (as indicated by the JDS equalling 1 in Figure 6a). However, when the adversary is unaware a stop has been made, this leads to more persistent divergences in the adversary observations and the true route as the adversary believes that stops have occurred at a different point in the route than actually occurred.

4.4 Data Privacy Techniques

As discussed in Section 2 observation privacy techniques rely on data privacy techniques being in use to address data privacy threats. For example, this might include encryption to ensure confidentiality of data in transit. These techniques are assumed to be used as a foundation which observation privacy techniques build upon. In this section, to compliment the representative techniques in each of the four observation privacy classes, we identify what data privacy techniques would need to be used in this system.

4.4.1 Protecting Data in Transit. The confidentiality of data in transit can be protected using protocols. In our transport example, communication between the smartcard and the reader can be encrypted. Similarly, communication between readers and backend systems can be protected using TLS. These techniques ensure that

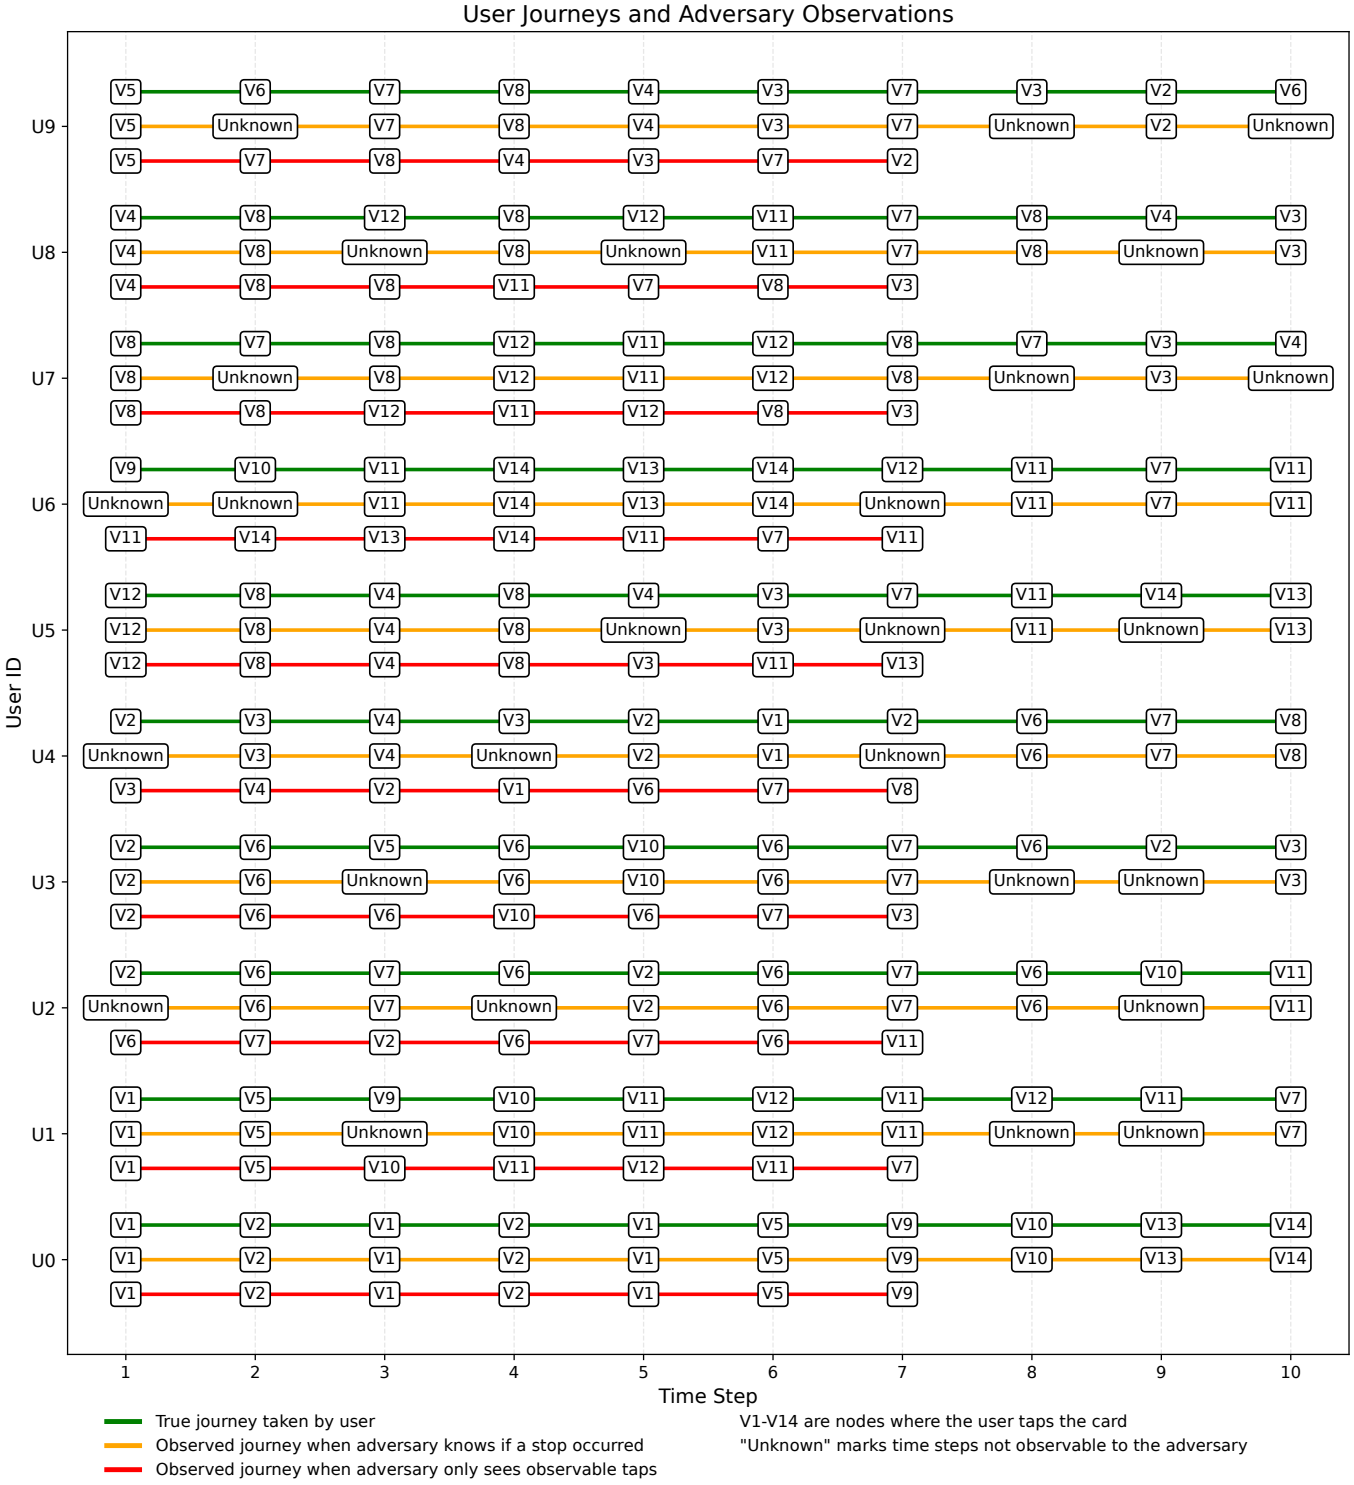

Figure 5: Adversarial observable paths and Actual Paths.

an attacker cannot read or modify the actual data being transmitted. For example, the smartcard identifier or transaction details are protected during communication. However, even with encryption, an attacker may still observe when communication happens, how often it occurs, or which devices are involved. This information can still reveal patterns of behaviour.

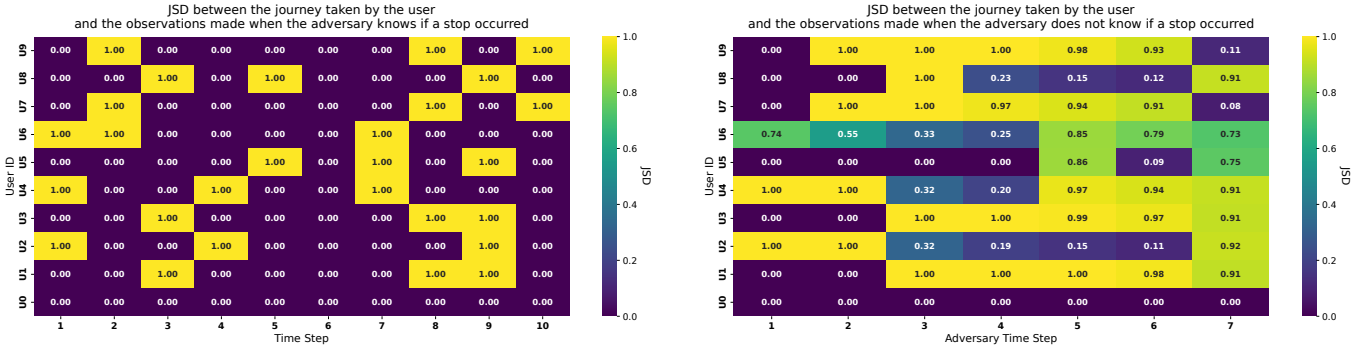

(a) The adversary knows a stop occurred.

(b) The adversary does not know a stop occurred.

Figure 6: Jensen-Shannon Divergence between the probability distributions of actual stops and adversary observations.

4.4.2 Protecting Data at Rest. Data stored in backend systems, such as travel history or transaction logs, can be protected using encryption and access control. This prevents unauthorised access to stored records. For example, databases storing smartcard usage can encrypt sensitive fields and restrict access to authorised entities only. Attacks such as privilege elevation or SQL injection have been widely studied and are out of scope of this work.

4.4.3 Private Computation over Data. Another desirable approach is undertaking a private computation over data without revealing it. This is commonly achieved using homomorphic encryption [25], however, these cryptographic operations are typically more expensive than traditional cryptography. Such protocols will still be vulnerable to observation threats to packet headers, timing of when messages are sent, direction of transmission (if wireless), and others that homomorphic encryption is not suitable to mitigate.

4.4.4 Private Release of Data. There may be a desire to allow third parties to query transport data and techniques such as differential privacy [16] can be used to bound the information revealed via these queries. These methods add controlled noise to datasets or aggregate data before release. For example, transport usage statistics can be shared without exposing individual travel patterns. However, an adversary directly observing the system does not rely on released datasets. Protecting released data does not stop inference from direct observations.

5 Related Work

Many existing studies address privacy leakage caused by observable system behaviour rather than direct data disclosure. In cyber-physical systems, an adversary can infer sensitive user context by observing locations, timing patterns, trajectories, or system interactions. This problem appears across different domains, including sensor networks, connected vehicles, smart homes, IoT platforms, UAVs, web systems, and via other side-channels [33]. Most works assume a passive or global observer who does not modify the system but collects repeated observations over time.

In this section, we review prior work using three classes of observation privacy techniques: (i) *Add Noise*, (ii) *Decorrelate*, and (iii) *Change Observability*. These classes help to explain how different systems reduce an adversary's ability to infer sensitive behaviour. Table 1 summarises how existing work fits into this classification. Our intent with this related work is to identify where observation privacy techniques have seen significant work and identify gaps in areas where certain observation privacy classes have not been used or domains in which observation privacy threats have not been extensively considered.

5.1 Add Noise

Multiple different areas have used techniques which add noise to the system behaviour to obscure relevant sensitive observations. In wireless sensor networks, a local observing adversary watches the direction from which messages originate and by following these over multiple hops can discover the location of a valuable asset being monitored. The seminal work introduced Phantom Routing, which sends messages through random intermediate nodes before reaching the final destination [24] which increases the amount of time to find the asset. However, this comes at the cost of increased latency. In contrast, fake source techniques introduce additional sensor nodes that act as decoy sources [5]. These approaches actively mislead a global observer by creating multiple plausible locations for the asset, however, this comes with a significant energy cost due to the high network traffic. Both approaches add noise, but they differ in where the uncertainty is introduced.

Connected vehicles may experience similar wireless eavesdropping to wireless sensor networks, however, for this system the adversary is typically attempting to track vehicles using digital identifiers over a long period of time to build a pattern of life of what vehicle. Noise can be injected for connected vehicles in different ways. One example is to obfuscate the information contained within the Cooperative Awareness Messages (CAMs) of two collaborating vehicles when their routes overlap [43]. Alternatively, vehicles can generate dummy queries [12] to a location based service. When adding noise this approach of generating fake events is a common approach used across multiple types of systems.

In UAV scenarios, an adversary may attempts to infer the drone's destination or mission intent from its trajectory. Noise can be introduced to its behaviour by altering its flight path through random intermediate locations [10, 13]. This is similar to Phantom Routing, but instead of network traffic it is the physical asset that changes

Table 1: Classification of Techniques used in Prior Work

Area	Paper	Technique Summary	i	ii	iii
Location Privacy in Sensor Networks	[24]	Add *phantom walks* to mislead adversary into believing message came from different location.	✓	-	-
	[30, 42]	When adversary has global visibility have all nodes send a message periodically/probabilistically to prevent adversary from locating the source.	-	✓	-
	[5]	Introduce fake sources, where sensor nodes act as a source to lure the adversary to their position instead of the genuine position of the asset.	✓	-	-
	[31]	Route messages to a data mule that physically transports (iii) the message to another part of the network (ii) before broadcasting it.	-	✓	✓
Connected Vehicles	[43]	Introduces obfuscation (i) to CAM messages to prevent a global passive eavesdropper linking successive messages (ii).	✓	✓	-
	[8, 9]	Vehicles change pseudonyms at same time and location (ii) during introduced silent periods where no messages are sent (iii).	-	✓	✓
	[12]	Relays dummy location-based service queries (i) through other vehicles (ii) to hide true locations and prevent deanonymisation under a global passive observer.	✓	✓	-
Smart Home Energy	[23, 41]	Use batteries to shape power consumption and hide the presence of specific appliances	-	✓	-
	[35]	Use batteries charged by renewable energy to supply power for the home during peak times (ii). Renewable generation alters what is visible at the smart meter as renewable load does not going through the smart meter (iii).	-	✓	✓
	[34]	Inject load signatures into a household's power consumption to emulate user presence (i) and hide device activity using solar panels (ii).	✓	✓	-
IoT	[40]	When interacting with IFTTT generate false events and include with true events when sending to IFTTT	✓	-	-
	[1, 36]	Injects network packets to prevent adversarial inference via eavesdropping.	✓	-	-
	[3]	Reduce a passive eavesdropper's ability to infer information from IoT network traffic by blocking traffic or using a VPN (iii) and shaping network traffic (ii).	-	✓	✓
UAV	[10]	Obfuscate the location and trajectory of a drone to a target from an observing adversary by visiting random intermediate locations along the route	✓	-	-
	[37]	Changes drone routes and pseudonyms within mix zones to break trajectory linkability	-	✓	-
	[13]	Prevent linking which vendor supplied an item to a destination by (ii) visiting all vendors before visiting all destinations and introducing (i) decoy vendors.	✓	✓	-
Additive Manufacture	[4]	Mitigate acoustic side-channel attack on additive manufacturing which aims to reproduce nozzle locations via introducing additional nozzle movements.	✓	-	-
	[2]	To make G-code reconstruction harder from acoustic side-channel, add (i) randomness to motor movements and (iii) soundproofing around the 3D printer to reduce the noise emitted.	✓	-	✓
Web	[14]	Onion Routing: Change the path that messages take through a network (ii) to prevent end-to-end tracing of a message via multiple layers of encryption (iii)	-	✓	✓
	[29]	TrackMeNot: Injects synthetic web search queries to hide genuine queries	✓	-	-

i Add Noise, ii Decorrelate, and iii Change Observability

the path being taken. Similarly, there is an increased journey length and energy use as a trade-off.

Smart homes will typically need to interact with services outside a home. Here, an adversary tries to infer user behaviour or device activity from network traffic patterns. The solution is to inject synthetic packets to protect against inference attacks based on traffic analysis [1, 36]. By injecting additional packets, the system increases network activity to make genuine communication harder to identify. In web services, similar ideas are applied at the application level. Here, the adversary observes search queries submitted by users and approaches, such as TrackMeNot, generate synthetic search queries alongside genuine user queries [29]. By mixing real

and fake queries, the system hides genuine queries along with artificial ones. The adversary cannot easily identify which queries reflect real intent.

In smart manufacturing, side-channels (from sources such as: acoustic [4], power [18], magnetic [22]) can reveal information, including details on what a 3D printer is manufacturing. Additional unnecessary movement can be added to the nozzle to obfuscate the shape that the 3D printer is producing [4].

Techniques which add noise to system behaviour introduce additional observations for the adversary which decreases certainty in what genuine actions have been observed. A high cost (e.g., energy, network traffic, latency) is typically incurred by the system to introduce these additional actions, so the techniques are only effective when systems can tolerate this additional overhead.

5.2 Decorrelate

Decorrelating techniques have been commonly used to change system behaviour by breaking a causal link or relationship between consecutive observations. Adversaries will typically use digital identifiers and timing of actions to link entities and infer behaviour, so these are what techniques focus on decorrelating.

To prevent linking one Cooperative Awareness Message (CAM) with a subsequent one, connected vehicles need to change the identity associated with this message. However, just changing the identity is insufficient as it can be re-linked. To address this, at locations where it is safe to do so (e.g., when stationary at an intersection), pseudonyms are changed while no CAMs are broadcasted [8, 9]. This silent period prevents the adversary linking between the old and new identity as long as there are a sufficiently large number of vehicles at the intersection. These locations for identity change are also used for UAV systems, where drone change routes or identities inside *mix zones* to prevent long-term trajectory tracking [37].

An alternate threat is where an adversary tries to associate repeated location-based service queries with a single requester. To mitigate this location-based service queries can be sent via others vehicles [12]. By forwarding queries through multiple vehicles, observations are distributed across entities. This weakens the link between a vehicle's query and the true requester.

In wireless sensor networks, techniques for a local passive observer are insufficient to address the threat of a global passive observer. So instead of perturbing routes, the detection of a valuable asset and the routing of that message can be decorrelated by having every node in the network broadcast a message periodically or probabilistically irrespective of if they have a genuine event to report [30, 42]. This comes with both a high energy cost and high latency to receive a message due to transmitters needing to wait until the next period before sending a message.

In smart home energy systems, an adversary attempts to infer occupancy or behaviour from power usage patterns. This observation is commonly performed via a smart meter. To decorrelate genuine activity from adversary observations energy storage (i.e., batteries) are commonly used to shape energy consumption to hide actives as power can be drawn from the energy storage instead of the grid [23, 41]. Renewable generation (e.g., solar panels) can also be used to shape power consumption from the grid which is visible to the adversary [34]. Alternatively, artificial load patterns are injected to simulate normal household activity and hide real appliance usage [34]. While these approaches weaken correlation across time, they incur additional energy costs or increase system complexity by introducing new devices.

Decorrelate techniques are also used in non-cyber physical systems. For example, Onion Routing uses layered encryption such that each network hop only knows the previous and next destinations in a route. By perturbing the route with multiple hops such as this end-to-end linkage of communication is prevented [14]. Similarly, in IoT settings, traffic shaping reduces a passive eavesdropper's ability to infer behaviour from network traffic patterns [3] or in UAVs by reducing an adversary's ability to link source and destination when delivering goods [13].

5.3 Change Observability

Change observability reduces privacy risk by limiting what the adversary can observe. In vehicular systems, ceasing to broadcast messages during pseudonym changes limits visibility [8]. As this only occurs while the vehicle is stationary, it is not an effective privacy preservation technique in isolation and requires the Decorrelation approach as previously discussed. However, the Decorrelation aspect of changing pseudonyms would not be effective without changing the observability of the system.

For Wireless Sensor Networks, changing the observability allows physical transfer of a message to be used to decorrelate routing of the message [31]. This assumes that an adversary only observes the wireless spectrum and does not have other observation capabilities (e.g., cameras) which could be used to link wireless communication to physical movement.

In smart home energy systems, renewable energy sources and batteries can supply power without being observed by the smart meter [35]. When adversary attempts to infer user behaviour from smart meter observations not all energy consumption is observable. In cases where not all the power consumption of a building is able to be met with renewables there will be some privacy loss.

In web systems, encryption is essential for data privacy, but using it against this threat model is also capable of mitigating some observation privacy threats. For example with Onion Routing [14] this prevents the entire path being known to every hop due to the multiple layers of encryption. In other settings VPNs and traffic blocking reduce the amount of metadata visible to passive observers [3].

6 Discussion

This work groups observation privacy techniques into three classes: (i) **Add Noise**, (ii) **Decorrelate**, and (iii) **Change Observability** which have been explained in Section 3. In this section, we discuss how they work in practice, their trade-offs, and how suitable they are for real-world cyber-physical systems. Each class reduces information leakage in a different way. However, how well they work depends on the system, the threat model, and system constraints.

Add Noise reduces the confidence of the adversary by adding extra or misleading behaviour. This makes it difficult to separate real actions from fake ones. However, this comes with cost. It can increase resource usage, delay system actions, and make behaviour harder to understand. Because of this, it may not be suitable for safety-critical systems.

Decorrelate techniques try to break the link between observations. This stops the adversary from connecting actions over time or between entities. These methods work well when there are many users or components. But if the system has only a few entities, the ability to provide privacy is limited.

Change Observability tries to reduce what the adversary can see. This can give strong privacy if it is possible. However, in practice, adversaries may observe the system in many ways [6]. Because of this, it is not always possible to fully hide system behaviour without affecting system performance or functionality. These classes are not used alone. In practice, systems may combine techniques from different classes to improve privacy. However, combining methods increases system complexity and overhead. So, the choice depends on system requirements and constraints.

6.1 Trade-offs Between Privacy and System Constraints

Cyber physical systems are often resource-constrained [21], which may be in terms of computational capabilities, memory (both stable and unstable storage), energy availability, communication limitations and often many other dimensions. Making use of observation privacy techniques will come with additional costs to cyber physical systems such as financial costs, higher latency, increased energy usage, reduced accuracy, additional computational and communication overhead, and others. Moreover, embedding privacy mechanisms can increase architectural complexity, which may make the system harder to operate, audit, or maintain, while at the same time concealing residual vulnerabilities. This overhead of observation privacy techniques can make it challenging to justify the introduction of these techniques.

Additionally, observation privacy techniques do not simply add potentially high overheads, but they may place demands on the system for the observation privacy technique to be feasible to implement and deploy. For example, our example of a *Decorrelate* technique draws inspiration from Mix Zones used in vehicular privacy techniques [8, 9, 37] which requires that there are a large number of entities (passengers, vehicles, UAVs, etc.) who are co-located before it is reasonable for identities to be swapped. The lower the number of entities in the system the easier an adversary can link the change in identity back to the original identities.

6.2 Speculative Class of Make Sensitive Commonplace

In addition to the three observation privacy classes observed in the literature across different cyber physical system domains, we also expect there to be techniques which can obscure underlying intent by making sensitive actions commonplace alongside common non-sensitive actions. This speculative class has been inspired by information theoretic privacy metrics which measure the surprise of the observing adversary [38]. The aim is to normalise the occurrence of sensitive behaviour and by doing so reduce the *surprise* of the adversary when it observes a sensitive event. The intuition is that events that are rare or unusual are of interest to an observer, by normalising these events, they cease to be interesting observations.

The adversary for this technique class is attempting to detect the occurrence of these sensitive events. For example, this could be to discover a change in the goal of a system by detecting actions that the system was not previously performing, but now needs to perform to achieve this new goal. For example, if an autonomous drone swarm is interested in surveilling a location that the swarm has not previously visited before, this is a surprising activity for an adversary to observe due to the deviation from standard behaviour. A technique in this class would be to incorporate this surveillance into usual behaviour even if the surveillance is not necessary. This normalises the occurrence of the sensitive event, and leads to less surprise by the adversary when observing the event. However, this is not a technique that appears to have been commonly used in literature and as such may not represent a useful classification of mitigations to an adversary who is directly observing the system.

6.2.1 Refined Threat Model. The adversary observes user journeys through smartcard tap in/out to look for unusual activity which differs to normal activity. Such anomalies highlight a change in behaviour which may be useful for an adversary to target their future actions.

6.2.2 Privacy Goal. In this scenario, the observer sees every destination the user visits and can recognise when a location is unusual to visit or is infrequently visited. Sensitive destinations stand out precisely because they are rare compared to routine movements. The user's goal is to decrease the perceived interestingness of these visits, while continuing to make them.

6.2.3 Example Technique. This technique incorporates sensitive activities into non-sensitive behaviour in order to normalise the sensitive actions. If an action, such as visiting a location, is rare and sensitive, then the user may perform this action when it is not needed within a journey that includes visits to non-sensitive locations. By including the unneeded visit to sensitive locations, the adversary is unable to determine when the user genuinely wanted to visit the sensitive location. Using this approach, the user is not trying to obscure the locations visited, but instead obscure that there was a motivation to visit the sensitive location.

To evaluate our illustrative technique, we assume a user takes a journey from V_1 to V_2 ($J_{\text{nonsensitive}} = V_1 \to V_2$) and does not mind if this pattern is observed. Occasionally, the user visits a sensitive location V_3 instead ($J_{\text{sensitive}} = V_1 \to V_3$), which they wish to conceal. When applying this class of technique to protect privacy, the user includes a visit to V_3 into all journeys, forming a combined journey $J_{\text{combined}} = V_1 \to V_3 \to V_2$. This makes visits to V_3 appear routine.

6.2.4 Information-Theoretic Analysis. Assume the following probability distribution $\mathcal{J}$ of which journey a user takes.

$$\Pr(\mathcal{J} = j) = \begin{cases} 0.8 & \text{if } j = V_1 \to V_2 \\ 0.2 & \text{if } j = V_1 \to V_3 \end{cases} \tag{1}$$

The self-information of an event x is a measurement of surprise in the occurrence of a random event [38], and is defined as:

$$I(x) = -\log_2 \Pr(x). \tag{2}$$

Rare events (like visiting V_3) have a high self-information (level of surprise) due to the low probability of the event occurring. By making V_3 part of every journey, its occurrence becomes less surprising due to the increase in probability of the adversary observing

the event occurring. When V_3 is part of every journey, its self-information drops to zero (as shown in Table 2).

Table 2: Effect of including sensitive location V_3 in all journeys.

Journey	Probability	Self-Information (bits)
Sensitive: $V_1 \to V_3$	0.2	2.32
Routine: $V_1 \to V_2$	0.8	0.32
After: $V_1 \to V_3 \to V_2$	1.0	0.00

6.3 Non-represented Areas

In this work we surveyed a range of different systems in Section 5, however, there are notable cyber physical systems that are absent from this section, such as: industrial control systems and general robotics applications. While this absence does not imply that observation privacy threats do not exist in these domains, it may mean that this kind of privacy problem is of little relevance to these areas.

One possible reason is that observation-based privacy threats are not traditionally viewed as a critical concern in these systems. Security research in such domains often focuses on safety, availability, and fault tolerance rather than inference attacks. As a result, privacy risks arising from system observability may be underexplored.

Another reason may be that these threats exist but have not yet been explicitly identified or labelled as observation privacy problems. For example, relevant studies may exist but were not captured due to differences in terminology or research focus. For some areas, the input which triggers system actions is the target of obfuscation. For example, with manufacturing there is a desire to obfuscate CAD models to mitigate counterfeiting [19]. However, this is a different threat model where the input to a system is being transformed, compared to the threat considered in this work where the behaviour of a system needs to be changed to mitigate an adversary directly observing the system.

Therefore, there may be scope for considering novel observation privacy threats to such systems, as part of eliminating reconnaissance steps in a Cyber Kill Chain.

7 Conclusion

In this work, we have introduced a classification of three observation privacy technique classes: (i) Add Noise, (ii) Decorrelate, and (ii) Change Observability, plus a speculative class of (iv) Make Sensitive Commonplace. The aim of which was to identify commonalities in observation privacy problems across a range of cyber physical system (and other) areas. We use a public transport scenario to illustrate what kind of technique might fall into each of these classes and then classified related work across six areas into these technique classes. This related work identified that adding noise to system behaviour and decorrelating the observations adversaries make and the states or actions of the system were the most prevalent approaches used to provide observation privacy. Fewer works changed the observability of the system to the adversary. This indicates that there could be scope for future work mitigating observational privacy attacks on cyber physical systems using this technique plus the speculative technique of making sensitive actions commonplace.

Additionally, some kinds of cyber physical systems (e.g., Industrial Control Systems) had no clear examples of work mitigating observational privacy threats. This could be due to the safety requirements of such systems making applying these technique classes infeasible, the lack of a suitable threat that necessitates applying these techniques, or potentially due to a different framing of the problem in the literature as different CPS areas use different terminology to refer to this problem.

Data Statement

The code used in this paper is available at: https://github.com/sampathkcs/Classes-of-Cyber-Physical-System-Observation-Privacy-Techniques.

Acknowledgments

This research was funded by the Engineering and Physical Sciences Research Council [EP/X040038/1].

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*. Association for Computing Machinery, New York, NY, USA, 207–218. doi:10.1145/3395351.3399421

[2] Mohammad Abdullah Al Faruque, Sujit Rokka Chhetri, Arquimedes Canedo, and Jiang Wan. 2016. Acoustic Side-Channel Attacks on Additive Manufacturing Systems. In *ACM/IEEE 7th International Conference on Cyber-Physical Systems (ICCPS)*. 1–10. doi:10.1109/ICCPS.2016.7479068

[3] Noah Apthorpe, Dillon Reisman, and Nick Feamster. 2017. POSTER: A Smart Home is No Castle: Privacy Vulnerabilities of Encrypted IoT Traffic. In *Proceedings of the 23rd Annual Network and Distributed System Security Symposium (NDSS)*. Internet Society, San Diego, California, USA.

[4] Seyed Ali Ghazi Asgar and Narasimha Reddy. 2026. QuietPrint: Protecting 3D Printers Against Acoustic Side-Channel Attacks. arXiv:2602.02198 [cs.CR] https://arxiv.org/abs/2602.02198

[5] Matthew Bradbury, Arshad Jhumka, and Matthew Leeke. 2018. Hybrid Online Protocols for Source Location Privacy in Wireless Sensor Networks. *J. Parallel and Distrib. Comput.* 115 (May 2018), 67–81. doi:10.1016/j.jpdc.2018.01.006

[6] Matthew Bradbury, Phillip Taylor, Ugur Ilker Atmaca, Carsten Maple, and Nathan Griffiths. 2020. Privacy Challenges with Protecting Live Vehicular Location Context. *IEEE Access* 8 (Nov. 2020), 207465–207484. doi:10.1109/ACCESS.2020.3038533

[7] Anas Bushnag, Abdelshakour Abuzneid, and Ausif Mahmood. 2018. Source anonymity against global adversary in wsns using dummy packet injections: A survey. *Electronics* 7, 10 (2018), 250. doi:10.3390/electronics7100250

[8] Levente Buttyán, Tamás Holczer, André Weimerskirch, and William Whyte. 2009. Slow: A practical pseudonym changing scheme for location privacy in vanets. In *2009 IEEE vehicular networking conference (VNC)*. IEEE, 1–8. doi:10.1109/VNC.2009.5416380

[9] Yi-Ming Chen and Yu-Chih Wei. 2013. A beacon-based trust management system for enhancing user centric location privacy in VANETs. *Journal of Communications and Networks* 15, 2 (2013), 153–163. doi:10.1109/JCN.2013.000028

[10] Samhith Reddy Chinthi-Reddy, Sunho Lim, Gyu Sang Choi, Jinseok Chae, and Cong Pu. 2022. DarkSky: Privacy-preserving target tracking strategies using a flying drone. *Vehicular Communications* 35 (2022), 100459. doi:10.1016/j.vehcom.2022.100459

[11] Mauro Conti, Jeroen Willemsen, and Bruno Crispo. 2013. Providing Source Location Privacy in Wireless Sensor Networks: A Survey. *IEEE Communications Surveys and Tutorials* 15, 3 (2013), 1238–1280. doi:10.1109/SURV.2013.011413.00118

[12] George Corser, Huirong Fu, Tao Shu, Patrick D'Errico, Warren Ma, Supeng Leng, and Ye Zhu. 2014. Privacy-by-Decoy: Protecting location privacy against collusion and deanonymization in vehicular location based services. In *2014 IEEE Intelligent Vehicles Symposium Proceedings*. IEEE, 1030–1036. doi:10.1109/IVS.2014.6856595

[13] Geoffrey Ding, Alex Berke, Karthik Gopalakrishnan, Kwassi H. Degue, Hamsa Balakrishnan, and Max Z. Li. 2022. Routing with Privacy for Drone Package

Delivery Systems. In *International Conference on Research in Air Transportation*. https://hdl.handle.net/1721.1/145404

[14] Roger Dingledine, Nick Mathewson, Paul F Syverson, et al. 2004. Tor: The second-generation onion router.. In *USENIX security symposium*, Vol. 4. 303–320. doi:10.21236/ADA465464

[15] 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, 265–284.

[16] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends® in Theoretical Computer Science* 9, 3–4 (2014), 211–407. doi:10.1561/0400000042

[17] Yebo Feng, Jun Li, Jelena Mirkovic, Cong Wu, Chong Wang, Hao Ren, Jiahua Xu, and Yang Liu. 2025. Unmasking the Internet: A Survey of Fine-Grained Network Traffic Analysis. *IEEE Communications Surveys & Tutorials* 27, 6 (2025), 3672–3709. doi:10.1109/COMST.2025.3545541

[18] Jacob Gatlin, Sofia Belikovetsky, Yuval Elovici, Anthony Skjellum, Joshua Lubell, Paul Witherell, and Mark Yampolskiy. 2021. Encryption is Futile: Reconstructing 3D-Printed Models Using the Power Side-Channel. In *Proceedings of the 24th International Symposium on Research in Attacks, Intrusions and Defenses* (San Sebastian, Spain) *(RAID '21)*. Association for Computing Machinery, New York, NY, USA, 135–147. doi:10.1145/3471621.3471850

[19] Nikhil Gupta, Fei Chen, Nektarios Georgios Tsoutsos, and Michail Maniatakos. 2017. ObfusCADe: Obfuscating Additive Manufacturing CAD Models Against Counterfeiting: Invited. In *Proceedings of the 54th Annual Design Automation Conference 2017* (Austin, TX, USA) *(DAC '17)*. Association for Computing Machinery, New York, NY, USA, Article 82, 6 pages. doi:10.1145/3061639.3079847

[20] Gerhard P. Hancke. 2011. Practical eavesdropping and skimming attacks on high-frequency RFID tokens. *Journal of Computer Security* 19, 2 (2011), 259–288. doi:10.3233/JCS-2010-0407

[21] Abdulmalik Humayed, Jingqiang Lin, Fengjun Li, and Bo Luo. 2017. Cyber-Physical Systems Security—A Survey. *IEEE Internet of Things Journal* 4, 6 (2017), 1802–1831. doi:10.1109/JIOT.2017.2703172

[22] Amirhossein Jamarani, Yazhou Tu, and Xiali Hei. 2025. Practitioner Paper: Decoding Intellectual Property: Acoustic and Magnetic Side-Channel Attack on a 3D Printer. In *Security and Privacy in Cyber-Physical Systems and Smart Vehicles*, Xiali Hei, Luis Garcia, Taegyu Kim, and Kyungtae Kim (Eds.). Springer Nature Switzerland, Cham, 54–74.

[23] Georgios Kalogridis, Costas Efthymiou, Stojan Z. Denic, Tim A. Lewis, and Rafael Cepeda. 2010. Privacy for Smart Meters: Towards Undetectable Appliance Load Signatures. In *2010 First IEEE International Conference on Smart Grid Communications*. IEEE, 232–237. doi:10.1109/SMARTGRID.2010.5622047

[24] Pandurang Kamat, Yanyong. Zhang, Wade Trappe, and Celal Ozturk. 2005. Enhancing Source-Location Privacy in Sensor Network Routing. In *25th IEEE International Conference on Distributed Computing Systems*. IEEE, Columbus, OH, USA, 599–608. doi:10.1109/ICDCS.2005.31

[25] Junsoo Kim, Chanhwa Lee, Hyungbo Shim, Jung Hee Cheon, Andrey Kim, Miran Kim, and Yongsoo Song. 2016. Encrypting Controller using Fully Homomorphic Encryption for Security of Cyber-Physical Systems. *IFAC-PapersOnLine* 49, 22 (2016), 175–180. doi:10.1016/j.ifacol.2016.10.392 6th IFAC Workshop on Distributed Estimation and Control in Networked Systems NECSYS 2016.

[26] Yuan Luo, Long Cheng, Hongxin Hu, Guojun Peng, and Danfeng Yao. 2021. Context-Rich Privacy Leakage Analysis Through Inferring Apps in Smart Home IoT. *IEEE Internet of Things Journal* 8, 4 (2021), 2736–2750. doi:10.1109/JIOT.2020.3019812

[27] Florence Mukamanzi, Manjula Raja, Tejodbhav Koduru, and Raja Datta. 2023. Position-Independent and Section-Based Source Location Privacy Protection in WSN. *IEEE Transactions on Industrial Informatics* 19, 5 (2023), 6636–6646. doi:10.1109/TII.2022.3183804

[28] Fabio Pasqualetti, Florian Dörfler, and Francesco Bullo. 2013. Attack Detection and Identification in Cyber-Physical Systems. *IEEE Trans. Automat. Control* 58, 11 (2013), 2715–2729. doi:10.1109/TAC.2013.2266831

[29] Sai Teja Peddinti and Nitesh Saxena. 2010. On the privacy of web search based on query obfuscation: a case study of TrackMeNot. In *Proceedings of the 10th International Conference on Privacy Enhancing Technologies* (Berlin, Germany) *(PETS'10)*. Springer-Verlag, Berlin, Heidelberg, 19––37.

[30] Alejandro Proano, Loukas Lazos, and Marwan Krunz. 2016. Traffic decorrelation techniques for countering a global eavesdropper in WSNs. *IEEE Transactions on Mobile Computing* 16, 3 (2016), 857–871. doi:10.1109/TMC.2016.2573304

[31] Mayank Raj, Na Li, Donggang Liu, Matthew Wright, and Sajal K. Das. 2014. Using data mules to preserve source location privacy in Wireless Sensor Networks. *Pervasive and Mobile Computing* 11 (2014), 244–260. doi:10.1016/j.pmcj.2012.10.002

[32] Nirupama Ravi, C. Mani Krishna, and Israel Koren. 2024. Mix-Zones as an Effective Privacy Enhancing Technique in Mobile and Vehicular Ad-hoc Networks. *ACM Comput. Surv.* 56, 12, Article 308 (Oct. 2024), 33 pages. doi:10.1145/3659576

[33] Vincenzo Rega, Luca Tari, Domenico Capriglione, Mario Molinara, and Fabrizio Marignetti. 2025. Side-Channel Measurements and Machine Learning for Classifying Application-Level Scenarios in IoT Contexts. In *IEEE International Workshop on Metrology for Industry 4.0 & IoT*. 123–128. doi:10.1109/MetroInd4.0IoT66048.2025.11122057

[34] Andreas Reinhardt, Dominik Egarter, Georgios Konstantinou, and Delphine Christin. 2015. Worried about privacy? Let your PV converter cover your electricity consumption fingerprints. In *2015 IEEE International Conference on Smart Grid Communications (SmartGridComm)*. 25–30. doi:10.1109/SmartGridComm.2015.7436271

[35] Jieqi Rong, Weirong Liu, Fu Jiang, Yijun Cheng, Heng Li, and Jun Peng. 2023. Privacy-aware optimal load scheduling for energy management system of smart home. *Sustainable Energy, Grids and Networks* 34 (2023), 101039. doi:10.1016/j.segan.2023.101039

[36] Akshaye Shenoi, Prasanna Karthik Vairam, Kanav Sabharwal, Jialin Li, and Dinil Mon Divakaran. 2023. iPET: privacy enhancing traffic perturbations for secure IoT communications. *Proceedings on Privacy Enhancing Technologies* 2023 (2023), 206–220. Issue 2. doi:10.56553/popets-2023-0048

[37] Alisson R. Svaigen, Azzedine Boukerche, Linnyer B. Ruiz, and Antonio A. F. Loureiro. 2021. MixDrones: A Mix Zones-based Location Privacy Protection Mechanism for the Internet of Drones. In *Proceedings of the 24th International ACM Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems* (Alicante, Spain) *(MSWiM '21)*. Association for Computing Machinery, New York, NY, USA, 181––188. doi:10.1145/3479239.3485712

[38] Isabel Wagner and David Eckhoff. 2018. Technical Privacy Metrics: A Systematic Survey. *ACM Comput. Surv.* 51, 3, Article 57 (June 2018), 38 pages. doi:10.1145/3168389

[39] Yonghui Xiao and Li Xiong. 2015. Protecting locations with differential privacy under temporal correlations. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security* (Denver, Colorado, USA) *(CCS '15)*. Association for Computing Machinery, New York, NY, USA, 1298–1309. doi:10.1145/2810103.2813640

[40] Rixin Xu, Qiang Zeng, Liehuang Zhu, Haotian Chi, Xiaojiang Du, and Mohsen Guizani. 2019. Privacy Leakage in Smart Homes and Its Mitigation: IFTTT as a Case Study. *IEEE Access* 7 (2019), 63457–63471. doi:10.1109/ACCESS.2019.2911202

[41] Weining Yang, Ninghui Li, Yuan Qi, Wahbeh Qardaji, Stephen McLaughlin, and Patrick McDaniel. 2012. Minimizing private data disclosures in the smart grid. 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, 415–427. doi:10.1145/2382196.2382242

[42] Yi Yang, Min Shao, Sencun Zhu, and Guohong Cao. 2013. Towards Statistically Strong Source Anonymity for Sensor Networks. *ACM Transactions on Sensor Networks* 9, 3, Article 34 (June 2013), 23 pages. doi:10.1145/2480730.2480737

[43] Yevhen Zolotavkin, Yurii Baryshev, Jannik Mähn, Vitalii Lukichov, and Stefan Köpsell. 2025. Optimal obfuscation of awareness messages: Improving users' unlinkability in Intelligent Transport Systems. *Computer Networks* 257 (2025), 110972. doi:10.1016/j.comnet.2024.110972