PRIVACY-PRESERVING CHALLENGES FOR URBAN SAFETY

Nesrine Kaaniche and Joaquin Garcia-Alfaro
Telecom SudParis, Institut Mines-Télécom & Institut Polytechnique de Paris[1]
kaaniche.nesrine@telecom-sudparis.eu, joaquin.garcia_alfaro@telecom-sudparis.eu

Abstract: We discuss privacy challenges for urban safety. We focus on privacy-preserving solutions that may allow urban safety operators to guarantee anonymous and unlinkable actions. We illustrate our analysis with the IMPETUS project scenarios and examine how design and architecture choices may impact compliance of personal data protection. Alternative designs that could lead to improvements in this matter are also briefly introduced.

Keywords: Privacy-preserving solutions, Privacy-Enhancing Technologies, Data Protection, Urban safety, Data Privacy Management.

1. Introduction

Multiple connected devices and sensors compose a smart city. These elements are connected through networks and their outputs are communicated to the inhabitants in an application via intelligent computing techniques. This relation between physical objects increases the possibility of turning cyberattacks into physical attacks. On one hand, if urban safety is not secured, this means that essential services may fail. However, too excessive security measures can also lead to other problems, such as privacy violations [1]. On the other hand, urban safety is supported by the massive collection of data. Thus, privacy and security have to be carefully addressed, finding a proper balance between them [2]. In this paper, we explore and discuss some challenges to the development and implementation of public urban safety with regard to technological solutions in smart cities. Our focus is on privacy-preserving solutions, whose goal is to allow urban safety operators to guarantee anonymous and unlinkable actions, while maintaining an appropriate degree of security. Through a case study based on the IMPETUS project [3], we examine some design and architecture choices, as well as how such decisions may affect compliance (e.g., technical and legal) of personal data protection. More specifically, we examine via the IMPETUS use case the problem of massive data collection, usually considered as the essence of emerging intelligent algorithms. Herein, we refer to Artificial Intelligence (AI). On one hand, the collection and manipulation of personal data raises alarming privacy issues, on the other hand, the learning algorithms, especially the most powerful ones, result in decision-making devices that are often not transparent and risk to be unfair. This refers to the privacy vs utility trade-off that has to be managed through the whole data life cycle. We will focus on the different Privacy Enhancing Technologies (PET) to enforce the privacy by design principle. We examine the effectiveness of such mechanisms that has been studied and demonstrated by researchers and with various pilot implementations. However, they are still not well perceived by many operators mainly because they are reputed to have an impact on the utility

[1] 19 place Marguerite Perey, 91120 Palaiseau, France.

This paper contributes to bridging the gap between the legal framework and the available technological implementation measures by providing an inventory of existing approaches and privacy design strategies of various degrees of maturity from research and development, from a service provider side. Starting

This paper is supported by European Union's Horizon 2020 research and innovation programme under grant agreement No. 883286, project IMPETUS (Intelligent Management of Processes, Ethics and Technology for Urban Safety).

from the privacy principles of the legislation, important elements are presented as a first step towards a design process for privacy-friendly systems and services. We also discuss alternative designs that could lead to privacy improvements. We explore how the use of Attribute based primitives can lead to the implementation of an anonymous certification framework, thus enhancing the data minimisation principle. The remaining sections are structured as follows. Section 2 provides some background. Section 3 focuses on the problem of massive data collection. Section 4 provides technical and legal solutions to the problem. Section 5 discusses alternative solutions based on ABS (Attribute Based Signature) for further privacy improvements. Section 6 concludes the paper.

2. Background

Connected smart cities are built upon a variety of services adapted to specific needs and citizens' expectations. Mainly designed with respect to personalization techniques, these services and applications rely on massive collection and analysis of gathered data. Indeed, a power imbalance between data processing entities, which determine what and how data is processed, and the individuals whose data is at stake, i.e., who might be influenced by decisions based on automated data analysis, or by failures to adequately protect private information, might be observed. To enforce privacy, the European Commission (EU) adopted, in 2018, the General Data Protection Regulation (GDPR) that sets up a legal context for the personal data collection, storage and processing. Privacy-Enhancing Technologies (PETs) [4, 5] have become a field that studies enabling techniques, investigates the level of data leakage, mitigates identification and traceability attacks and implements privacy-preserving processing. From this perspective, it is of utmost importance to first understand the main requirements and different security models to evaluate the relevant building blocks from the design phase. To develop different tools to be integrated in the wide city, while enhancing the privacy of citizens, we hereafter present the generic functional architecture and main interactions between data layers.

2.1. Functional architecture

For our analysis, we consider a functional architecture which consists of four layers with specific responsibilities and components. First, a *sensing layer* that consists of various equipment that collect data from the surrounding physical environment and share it with the data collection layer. For instance, different types of sensors, actuators or CCTV cameras are considered as main components of the sensing layer. Second, a *data collection layer*, that deals with the transmission of gathered data using reliable wired or wireless communication to the local or remote databases. This layer mainly involves the communication protocols and services, from a networking point of view. Let us emphasize that storing a big volume of data incurs high storage overhead on the existing databases of this layer. Thus, many applications would rely on remote edge or cloud servers to address this issue and remove the burden of maintaining large infrastructures, especially when data involves real-stream data. Referring to the GDPR nomenclature, these databases may be maintained by either data processors or data controllers.

The remaining two layers consist on *data processing* and *applications*. The data processing layer performs all "*pre*"-processing techniques, i.e., storing and analysing large amounts of data and maintaining infrastructures, with regards to different applications and services' requirements. The application layer is responsible for exchanging data between operators (e.g., citizens and stockholders) and smart applications. Exchanged data can be raw, aggregated or processed via accurate analysis and visualization algorithms.

2.2. Pillars infrastructure

In the context of IMPETUS' use cases, i.e., dealing with public security and safety, we identify four pillars' infrastructures, namely institutional, social, physical and economic, introduced as follows:

- *Institutional infrastructure:* deploys the fundamental activities, e.g., management, governance and planning of events. It relies on the citizens' collected and produced data in decision-making processes.

- *Physical infrastructure:* involves IT, communication and hardware components that yield to support a physical environment for urban safety and mobility. For instance, urban mobility focuses on the quality of cycling, pollution indicators, and smart transportation systems in cities, while public safety considers safe walking and secure gathering places and holding social events.
- *Social infrastructure:* involves diverse mechanisms to promote and develop human and social capital, and provide intelligent and straightforward connected infrastructure for addressing different social needs and services of citizens, such as environment, and inclusive planning.

- *Economic infrastructure:* refers to the basic services that help to promote the process of production and distribution of economic activities and develop proper infrastructure to generate employment opportunities and attract investments. Although this type of activity may not directly be represented by IMPETUS outputs, the different tools may have an impact on the economic development and the attractiveness of a city.

2.3. Security and privacy challenges in the context of the IMPETUS project

The IMPETUS project [3] has conducted a study with different focus groups whereas challenges related to privacy in a large connected city have been identified. Below, we organized a summary of the different raised challenges.

- Maintaining privacy during regular Internet browsing in connected cities is considered as "*difficult*". The first main concern is the widespread deployment of artificially intelligent processing algorithms that can be used in combination with the collected personal information to deduce involuntary correlations, leading to specific identification (other people, web pages, organizations, *etc.*).
- The second privacy concern is related to the tracking of spatial mobility, e.g., in relation to pedestrians, consumers and vehicles. Tracking is already a legitimate part of smart city technologies, as per ensuring safety in the public space, but there is a fear of misuse, e.g., related to unwanted surveillance.
- The third reported challenge consists of properly informing citizens about what the information is being used for, obtaining and maintaining informed consent in a practical manner. This challenge becomes even harder when combining different data sources. The focus groups expressed that aggregation of data may lead to profiling, discrimination, and political manipulation.

In order to understand the aforementioned challenges, we hereafter present threat models and main challenges, namely related to extracting information about data owners, and also referred to as citizens and users, or organizations, service providers and governmental institutions. Taking the example of smart mobility, it is imperative that not only the privacy of the collected and analysed data be preserved but also the running algorithms (usually considered as sensitive and proprietary). Regardless of the goal, the attacks and defences relate to exposing or preventing the exposure of analysis algorithms (processes) and collected data.

2.4. Threat models

Privacy and security risks are mainly related to environments, technologies and involved parties. Indeed, as pointed out in a recent report of the European Network and Information Security Agency (ENISA) [4, 5], understanding privacy concerns from a technical point of a view, leads to identify:

1. Collected and processed data that are released and may be considered as sensitive, personal and identifying data,
2. Data that may be used to identify and/or revoke the anonymity of a user,
3. Potential adversaries (i.e., actors that may gain access to personal identifying information) which can rely on data being transferred and processed that the adversary has access to, and external and background knowledge of the adversary- possible collusion with other entities.

   Recall that adversaries may be passive or active, and considered under either semi-trusted or untrusted environments [4], presented as follows:
   - Passive attacks: the adversary passively observes the data and performs inference or concludes connections, e.g., without changing anything in the process. These attacks are usually considered against privacy requirements, *e.g., anonymity, unlinkability, unobservability, etc.*
   - Active attacks: the adversary actively changes the data or processes. These attacks are usually considered against security requirements, *e.g., integrity, availability, etc.*

3. Massive data collection

As discussed in Section 2, IMPETUS identified several challenges with respect to the enforcement of privacy technical requirements and the compliance with the legislation. These challenges are very tight to technologies, the functional architecture and involved actors. In the following, we point out the different vulnerabilities and issues that may arise from a technologies' point of view, while referring to the architectural layers (cf. Table 1).

3.1. IoT-based challenges

The Internet of Things (IoT) consists of interrelated, internet-connected (smart) objects that are able to collect and transfer data over a wireless network without human intervention. There are various privacy issues associated with smart devices that are mainly due to the massive collection of data, focused on the sensing and communications layers. Indeed, the connected devices have the capability to be used as a mediator storage or a fog node to perform a small computation in the network. These sensing and preprocessing capabilities make them vulnerable end-points for collecting the exchanged data and enriching adversarial databases, thus conducting specific correlation and inference attacks. While a huge number of applications are continuously proposed to provide various benefits for citizens, the majority of these applications gain access to private information of users -*without acquiring explicit informed consent*- and may transfer the collected data to *unauthorized* third parties. Finally, the sensing capabilities of the smart devices *facilitate* the bypass of the data minimization principle, and most applications usually collect more data than the necessities of original functions, while in permission scope, which is known as data over-collection.

3.2. Cloud based issues

In order to cope with the shortcomings of smart devices, i.e., processing and storage capacities, battery constraints, etc. various applications delegate the data and processing management to *external* cloud providers. While outsourcing data and processing has various economic advantages, several security and privacy challenges are identified in [3]. Next, we summarize common challenges raised by cloud infrastructures, platforms and applications.

Table 1. Summary of privacy challenges with respect to IMPETUS functional architecture layers

	IoT	*Cloud*	*AI*
Sensing layer	Data over collection		- Data over collection - Data poisoning - Backdoor injection
Collection layer	Lack of standardized secure short-band communication protocols		
Processing layer	Limited computation resources for advanced secure (cryptographic) algorithms	-Loss of data and computation control -Lack of knowledge about effective SLA enforcement -Multi-tenancy	- Inference attacks - Model theft
Application layer	Open and insecure APIs		

- *Data and computation outsourcing:* by outsourcing the data to remote servers, data management is delegated to a third-party provider, usually considered as a semi-trusted or honest-but-curious entity. This raises privacy concerns, such as the anonymity of data owners.
- *Physical location of data:* the lack of knowledge about the physical location of data in cloud services may have an impact on the data security, quality of services and might harm users' privacy. This latter is of utmost importance as data legislation regarding the collection and processing of data is different between different countries and regions, and can be more intrusive compared to the EU regulations.
- *Lack of knowledge about Service Level Agreements (SLAs):* SLA is a contract signed between the client and the service provider including functional and non-functional requirements. It considers obligations, service pricing, and penalties in case of agreement violations. However, due to the abstract nature of clouds, SLA violations with regards to data involve data retention, privacy leakage.
- *Multitenancy:* this cloud feature means that the cloud infrastructure is shared and used by multiple users. In a nutshell, data belonging to different users may be located on the same physical machine, based on a specific resource allocation policy. Due to the multi-tenancy's economic efficiency, providers usually select this feature as an essential block for the cloud environment design. However, it generates new threats, such that, malicious users may exploit this co-residence issue to perform privacy (inference) attacks.

3.3. AI-based attacks

Recent progress in Artificial Intelligence (AI) in general, and Machine Learning (ML) in particular, is continuously encouraging many sectors to integrate AI-based algorithms in different processes. AI is a key enabler of smart cities, where the size and complexity of smart cities' systems are key challenges. The ability to efficiently and process gathered data and monitor in real time the state of critical infrastructures increasingly become an added value and a practical need. Unfortunately, they are

generally considered as data-hungry tools and their benefits are often accompanied by a mostly blackbox character and high complexity of the final algorithms in use, rendering conventional methods for safety assurance insufficient or inapplicable. Hence, the need to enforce privacy by design [6, 7].

As presented above, the massive collection of data from the different devices (i.e., when referring to the sensing and data collection layers), constitute first threat vectors to attack intelligent systems due to their multitude and their limitations in terms of resources and security features. For instance, by poisoning smart city's data, adversaries can try to fake the models, implying they will learn the correct correlation between data and the state of a critical system (modifying the model boundaries), or they can push the model in taking decisions that are hampering the city's infrastructure and population. In this context, there is the need to confirm the common assumption on the effectiveness of the employed ML models, adopting suitable privacy and security techniques. These techniques aim to counteract adversaries trying to deceive ML models at different layers: i) data collection and ingestion, ii) training, iii) inference. The state of the art in the domain shows that while the performance of ML and neural network architectures had a boost in the last years, their robustness to adversarial settings is asking a step ahead [8, 9, 10,11]. It is clear that the strong link between sensor data and ML models, as well as the intrinsic weakness of the sensors themselves, introduce new serious risks, introduced as follows:

- *Backdoor injection:* the first one aims to manipulate data to attack the learning phase. In this case, the attacker crafts and distributes corrupted data, which are used by ML algorithms to build an inaccurate model of the system behaviour. These attacks are referred to as Machine Learning Poisoning [13]. They produce a poisoned ML model that learns a wrong correlation between data and the state of the monitored infrastructure. Smart attacks based on careful data manipulation can open the door to stealth attacks on the infrastructure, providing adversaries with the ability to introduce *"backdoors"* in the model. These backdoors induce erroneous classification of inputs, with possibly disastrous consequences on the working of the whole system. For instance, if an anomalous detection model is trained, the machine learning poisoning can introduce a backdoor that impedes the model to classify an anomaly, identifying it as a safe behaviour. Let us consider a ML model that is monitoring the quality of air pollution. An adversary can fake the model in believing that the presence of some chemicals in the air is innocuous, while they could be dangerous for the population.
- *Data Poisoning:* the second one aims to inject specific data, generally carefully selected, to fake an existing model into taking decisions that decrease performance and increase risks of the city's infrastructure and population. This category of attacks, called adversarial examples [8], builds on sensor inputs that can trick a deployed model, trained on benign data, into making a wrong decision. The most difficult aspect here is that adversarial example attacks are difficult to counteract since poisoned data are generally indistinguishable from a normal input for humans. For instance, let us consider a monitoring service in a smart city, via CCTV cameras. Adversarial examples can be used by an attacker interacting with different cameras to let the model believe that certain areas of the city are congested. This would force the model to reroute the traffic towards busy areas, as well as changing the traffic light timing, creating a gridlock. This would have disastrous consequences for instance in the case of a terrorist attack.
- *Model theft:* the third one is a mix of the previous two and is employed in scenarios, and mainly considered a security threat where ML models are either i) retrained over time or ii) alternative models have been trained and can be deployed on the basis of contextual information.
- *Inference attacks:* the fourth category of attacks involves two main attacks, namely (a) inference about members of the population and (b) inference about members in the training set. For the first case (a), an adversary can use the model's output to infer the values of sensitive attributes used as input to the model. Note that it may not be possible to prevent this if the model is based on statistical facts about the population: for example, suppose that training the model has uncovered a high correlation between a person's externally observable phenotype features and their genetic predisposition to a certain disease; this correlation is now a publicly known fact

that allows anyone to infer information about the person's genome after observing that person. For the second case (b), the focus is on the privacy of the individuals whose data was used to train the model. For instance, given a model and an exact data point, the adversary infers whether this point was used to train the model or not. The adversary may also try to extract properties. In fact, training data may not be identically distributed across different users whose records are in the training set; unlike model inversion, the adversary tries to infer properties that are true of a subset of the training inputs but not of the class as a whole.

4. Privacy-preserving solutions

This section gives an overview of privacy enabling techniques and details their suitability to different smart city scenarios and identified challenges. It classifies PETs into two groups, namely server-side and channel-side solutions. Then, it gives an overview of existing legal solutions.

4.1. Technical solutions

4.1.1. Server-side solutions

Server-side enabling technologies include three main categories, i.e., data perturbation, secure processing and database anonymization.

4.1.1.1. Data perturbation

Data Perturbation aims at intentionally making information difficult to understand or perceive for security and privacy reasons. In fact, the speed of dissemination of information, the technical progress and the global nature of the Internet make it difficult to delete data that may be too personal, embarrassing or confidential. Thus, perturbation consists mainly in publishing large amounts of information that are false, imprecise, irrelevant and/or organized in such a way that the information that one wishes to protect is hidden, i.e., embedded in a large volume of data. Data perturbation techniques are used for enhancing privacy in various querying services. In order to protect queries, one idea consists of generating dummy queries that will be sent to the central server along with the real query. The main issues of these techniques are the privacy utility trade-offs induced by the suppression technique, removing some records or details [4, 5].

4.1.1.2. Secure Multiparty Computation

Privacy preserving computation techniques aim at protecting users' privacy and the secrecy of data contents during processing over these data. The goal of Secure Multiparty Computation (SMC) techniques is to enable distributed computing tasks among participating entities in a secure manner. That is, SMC considers that a group of participants wants to carry out a joint computation of a given function while keeping secret the input data of each party. SMC has been used to solve several privacy-preserving problems such as private database queries, secret voting, privacy preserving data mining and privacy preserving intrusion detection tools and mechanisms. Three different approaches are generally deployed to provide secure multiparty computation functionalities, namely oblivious transfer, homomorphic encryption, and secret sharing techniques. The oblivious transfer protocol generates high processing and communication overheads. The secret sharing approach gives better results in terms of computation cost, thanks to the usage of primitive operations. However, it requires the existence of secure channels between different participating entities, hence generating a high bandwidth consumption, due to the involved interactions between users. The homomorphic encryption does not require the existence of secure channels and assures a high level of privacy. However, it necessitates several processing operations to ensure homomorphism properties, thus generating high computation complexity.

4.1.1.3. Database anonymization

Database anonymization techniques are basically used to protect data within statistical databases. They permit to resolve the trade-off between data usability and users' privacy preservation, as revealed results, either the databases or a specific result over the database do not permit to reveal information related to a specific user. These techniques also include Differential Privacy mechanisms. Anonymization techniques are relevant for various use-cases, namely applications that do not require to learn the original user's identity, but only context information. Anonymization techniques mainly refer to database privacy preservation. Even so, for cooperative applications where the database belongs to several corporations, it comes to the privacy protection of the various collaborating entities.

Main techniques for anonymizing databases w.r.t. respondent, owner and users' privacy include *kanonymity, t-closeness* and *l-diversity* [4, 5]. Note that these techniques that are originally used over statistical databases have extended usage to dynamic data. Briefly, to implement k-anonymity, it is important to recognize which attributes are considered as key credits, called likewise semi identifiers. In other words, k-anonymity can forestall character divulgence, i.e., a record in the k-anonymized set S, k cannot be planned back to the comparing record in the first S, subsequently, by guaranteeing that each record is indistinguishable by essentially other k − 1 records dependent on the worth of key credits. For Location-Based Services (LBS), an attacker, having access to users' location, may be able to identify the requesting user, relying on its spatio-temporal parameters. Consequently, several research works propose to expand the precise location of the user to involve several potential requesting issuers. This leads to generalizing several context-data to ensure anonymity, thus resulting in the context information released to the service provider being sometimes too large and imprecise to provide an acceptable quality for the service.

Differential Protection (DP) is acquiring a growing interest, primarily to guarantee security saving information mining. In a nutshell, differential privacy ensures that the removal or addition of a single database item does not (substantially) affect the outcome of any analysis (i.e., the probability distribution of released items does not significantly change). This property is enforced by adding random noise to the exact outcome. Note that differential privacy addresses data leakage attacks as even if a user has removed his data from the data set, no outputs would become significantly more or less likely. When DP techniques are applied at the data owner side, without the need for a third party, it is called Local Differential Privacy (LDP), whose main idea is to allow users to locally perturb their input data.

4.1.2. Channel-side solutions

To secure communications against pervasive surveillance, several service providers propose to deploy encrypted communication channels. It is important to emphasize that encrypted channels need to be implemented and configured correctly, to ensure a sufficient security level. Several technologies and protocols have been introduced, namely the well-known Transport Layer Security 1.2 protocol (i.e., TLS 1.2) and the Secure Shell (SSH) protocols). These technologies provide a confidential and conceivably authenticated channel between users and service providers.

4.2. Legal and operational solutions

In an effort to provide an analysis of various aspects of privacy, security, and surveillance concerning the involuntary visual and audio capture of personal property, access to personal data, involuntary surveillance, storage, and security of data collected in a smart city context, this section recalls data privacy legislation with a focus on GDPR.

In 2018, the GDPR came into force for effectively ensuring the protection of the data subject's personal data. In particular, the regulation clarifies the conditions under which it is compulsory to obtain the consent of the data subject before processing his personal data, especially for sensitive personal data and

data relating to minors. The GDPR also introduces the new obligation of accountability for organizations (i.e., data processors and data controllers). Indeed, each entity processing personal data must be able to demonstrate at any time that it is complying with the obligations laid down by the GDPR. According to the GDPR, the data subject's consent is given for specific purposes that must be compliant to both the data controller and the data processor. In this context, three main roles are defined. The data subject who gives his consent to a data controller (i.e., organization,) for the processing of his personal data, with the possibility to forward them to a data processor (i.e., organization) that may process data on behalf of the data controller. Here, data controllers are responsible for:

1. specifying to the data subject the purpose of data collection,
2. obtaining the data subject's consent, and
3. processing personal data according to the consented purposes, and not beyond. We note that for ease of presentation, the remainder of the paper refers to the data subject as the data owner and to both the data controller as well as the data processor as the service provider.

From a data owner perspective, there is a need for new security mechanisms that support data accountability and provenance auditing. In a nutshell, these solutions have to ensure that personal data were accessed by data controllers and/or forwarded to data processors. Indeed, it is important to conceive a secure and transparent solution that permits data owners to (i) check that data controllers and processors are correctly using their personal data with respect to the consented purposes, (ii) verify whether data were accessed, processed, or forwarded without their consent, and (iii) withdraw their consent.

From a data controller or processor perspective, there is a need to design a trusted and transparent accountability solution that enables them to get a proof of the data owner's given consent prior to gathering, accessing, processing, or storing his personal data.

The NIST-PF is a voluntary tool developed in collaboration with stakeholders intended to help organizations identify and manage privacy risk to build innovative products and services while protecting users' privacy. Referring to NIST-PF, the privacy properties that have to be enforced by the designed services, are summarized as follows:

- Anonymity, i.e., the ability of the user to access a resource or service without disclosing his identity to third parties. That is, the anonymity of a user means that he is not identifiable within a set of subjects, known as the anonymity set. Several levels of anonymity have been defined in the literature, ranging from complete anonymity (i.e., no one can reveal the identity of the user) to pseudo- anonymity (i.e., the identity is generally not known but can be disclosed if necessary) to pseudonymity (i.e., multiple virtual identities can be created and used in different settings).
- *Data minimization,* a fundamental feature of privacy preservation, requires that service providers collect and process the minimum amount of information needed for appropriate execution of a service or a particular transaction. The goal is to minimize the amount of collected personal information by service providers, for instance, to reduce the risk of profiling and tracking users.
- *Unlinkability,* closely related to the anonymity property, refers to Items of Interest (IoIs, e.g., users, messages, actions, etc.) that, from an attacker's perspective, it is unfeasible to distinguish whether they are related or not.
- *Unobservability* refers to the difficulty of identifying a given user out of several other users involved in an IOI, as well as the anonymity of the user against the other users involved in a given IoI. In other words, unobservability refers to situations in which a user can use a resource or a service, without being noticed by others. It also requires that third parties cannot determine if an operation is running.

4.3. Discussion

Table 2 presents an overview of commonly deployed privacy mechanisms with respect to different enabling technologies. It also describes different implemented solutions relying on the underlying PET. Table 1 shows that PET are deployed in different contexts and various purposes, thus enabling privacy by design for the whole data lifecycle, from the sensing to the processing and application levels. However, it is important to emphasize that in order to ensure privacy properties and guarantee an acceptable level of privacy in a connected city, usually considered as an open and complex environment, it is crucial to consider the interactions between all different actors. That is to say, it might be insufficient to implement PET to ensure the privacy requirement of a specific service or application, i.e., as a standalone tool with no consideration of the surrounding environment.

Table 2. Commonly deployed privacy mechanisms with respect to different enabling technologies.

PET		**Technology**	**Description**
Secure Communications			Secure public WiFi with WPA2
		Ubiquitous Connectivity	Use anonymous communication to protect metadata, i.e., Tor
			Ensure correct usage of SSL/TLS with static analysis
Secure Multiparty Computation	**Private Information Retrieval**	Cloud Computing	Process data with private inputs, e.g., genomic tests
			Hide access patterns to remote files and databases
	Homomorphic Encryption		Perform privacy-preserving data mining over distributed datasets Privately process data at third parties
		Internet of Things	Aggregate data over multiple users
Database Anonymization		Internet of Things	Ensure *k*-anonymity of sensor readings
			Use *l*-diversity or hierarchical map quantization to prevent location inference attacks against *k*-anonymity

		Cluster IoT data streams and only release clusters with at *at least k*-members
	Open Data	Release only data that satisfy *k*anonymity, *l*-diversity, *m*invariance, and *t*-closeness
	Ubiquitous Connectivity	Change device identifiers frequently to prevent fingerprinting, randomize browser fingerprints, insert cover traffic
Differential Privacy	Internet of Things	Apply noise to meter readings
	Open Data	Release noisy aggregates of data, e.g., public transport data or t-closeness

5. Alternative solutions

Privacy preserving authentication, likewise known by privacy preserving certification or Attribute based certification (AC), are cryptographic systems that enable clients to acquire certified credentials associated with their attributes from trusted authorities, and later derive presentation tokens that reveal just required data fulfilling service providers (SP)' predicates. Diverse substantial developments have been proposed and considered as fundamental building blocks in identity management systems. Indeed, depending on AC, every client can demonstrate to a service provider that he holds validated properties, referred to likewise as credentials, obtained from issuing authorities. In addition, AC techniques prevent SPs to trace users' activities based on successive authentication sessions. That is, the user derives a proof associated to each different access request, such that the SP is not able to link a single received proof to another or to any information relative to its owner, even in case of collusion between providers and with the credential issuer.

AC methods attract a lot of interest and complete consideration from industries and academia, thanks to their capacity to help the data minimization basic component. The design of a privacy preserving certification scheme strongly relies on the use of malleable signature schemes that provide several interesting properties, such as the selective disclosure feature and the unforgeability property. In fact, the selective disclosure property refers to the ability provided to the user to present to the verifier partial information extracted or derived from his credential, for instance, to prove he is older than 18 to purchase liquors, while not revealing his birth date. The unforgeability property ensures that unless a user possesses a legitimate and certified credential, i.e., secret key, he is not able to generate a valid authentication proof, i.e., user's signature over the SP's access policy.

According to this viewpoint, malleable signatures are considered the key building blocks to build privacy preserving yet authenticated access, specifically ABS [12, 13], sanitisable signatures and group signatures [5], supporting the data minimization guideline. For example, ABS enables a client to sign a message with respect to a particular access structure defined over attributes. Every client, holding a bunch of properties, needs to acquire a private key related to his attributes from an issuing entity. Accordingly, the client can sign a message with regard to a predicate fulfilled by any subset of his certified attributes . The verifier cannot deduce more than the correctness of received signature, i.e., the client cannot then guess which attributes have been used.

6. Conclusion

This paper provides a review of most commonly deployed PET in the context of smart cities applications. First, it is important to emphasize that due to the diversity of smart applications, different privacy technologies need to be combined to ensure an acceptable level of privacy. Indeed, smart cities combine so many technological components that it is not enough to simply apply privacy technologies to each component. Instead, we advise that the interactions between technologies and data have to be considered to design *"joint privacy technologies."* This is especially important because applications start with isolated solutions that get integrated gradually. Thus, one approach to facilitate joint privacy protection is to focus on the interfaces between different systems, on their interactions and in particular on the data flow. For example, different components in a sensor-based application may all deploy independent differential privacy mechanisms before transferring data to the processing layer. Taking this into consideration will help to define appropriate privacy enabling mechanisms for the data storage and processing.

Acknowledgments

We acknowledge financial support from the European Commission (H2020 IMPETUS project, under grant agreement 883286).

References

[1] H. Habibzadeha, B. H. Nussbaumb, F. Anjomshoac, B. Kantarcid, and T. Soyat (2019). A survey on cybersecurity, data privacy, and policy issues in cyber-physical system deployments in smart cities. Sustainable Cities and Society, vol. 50.

[2] I. Agadakos, P. Anantharaman, G. F. Ciocarlie, B. Copos, M. Emmi, T. Lepoint, U. Lindqvist, M. Locasto, L. Song (2020). Chapter 9. Securing Smart Cities. Implications and Challenges. Modeling and Design of Secure Internet of Things, 1st Edition. John Wiley & Sons.

[3] IMPETUS project. https://www.impetus-project.eu/

[4] N. Kaaniche, M. Laurent, S. Belguith (2020). Privacy enhancing technologies for solving the privacy-personalization paradox: Taxonomy and survey. Journal of Network and Computer Applications.

[5] D. L. Chaum (1981). Untraceable electronic mail, return addresses, and digital pseudonyms. Communications of the ACM, 24(2), pp.84-90.

[6] L. Cui, G. Xie, Y. Qu, L. Gao and Y. Yang (2018). Security and Privacy in Smart Cities: Challenges and Opportunities. IEEE Access, vol. 6, pp. 46134-46145.

[7] A. Gharaibeh, M. Salahuddin, S. Hussini, A. Khreishah, I. Khalil, M. Guizani, A. Al-Fuqaha (2017). Smart Cities: A Survey on Data Management, Security, and Enabling Technologies. IEEE Communications Surveys & Tutorials, vol. 19, no. 4, pp. 2456-2501.

[8] I. Goodfellow, J. Shlens, C. Szegedy (2014). Explaining and harnessing adversarial examples. arXiv.

[9] IoT Security Foundation. Physical security best practices. Available On-line (Last Access: October 2021): https://www.iotsecurityfoundation.org/best-practice-guide-articles/physical-security/

[10] N. Papernot, P. McDaniel, S. Jha, M. Fredrikson, Z. B. Celik, A. Swami (2016). The limitations of deep learning in adversarial settings. 2016 IEEE European Symposium on Security and Privacy.

[11] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, R. Fergus (2013). Intriguing properties of neural networks. arXiv.

[12] N. Kaaniche, M. Laurent (2016). Attribute-based signatures for supporting anonymous certification. European symposium on research in computer security, pp. 279-300.

[13] N. Kiennert, N. Kaaniche, M. Laurent, P. Rocher, J. Garcia-Alfaro (2017). Anonymous certification for an e-assessment framework. In Nordic conference on secure IT systems, pp. 70-85.