Smart Meter Data Privacy: A Survey

Muhammad Rizwan Asghar, György Dán, Daniele Miorandi and Imrich Chlamtac

***Abstract*—Automated and smart meters are devices that are able to monitor the energy consumption of electricity consumers in near real-time. They are considered key technological enablers of the smart grid, as the real-time consumption data that they can collect could enable new sophisticated billing schemes, could facilitate more efficient power distribution system operation and could give rise to a variety of value-added services. At the same time, the energy consumption data that the meters collect are sensitive consumer information; thus, privacy is a key concern and is a major inhibitor of real-time data collection in practice. In this article, we review the different uses of metering data in the smart grid and the related privacy legislation. We then provide a structured overview, shortcomings, recommendations and research directions of security solutions that are needed for privacy-preserving meter data delivery and management. We finally survey recent work on privacy-preserving technologies for meter data collection for the three application areas: billing, operations and value-added services including demand response.**

***Index Terms*—Smart Grids, Smart Meters, Privacy, Cybersecurity**

I. Introduction

The term smart grid is used broadly to refer to the next generation of electrical energy transmission and distribution infrastructures, which will be characterised by a tight integration with Information and Communication Technologies (ICT). The integration of the power grid with ICT will enable pervasive real-time monitoring of the physical processes, including generation and consumption at the customers' premises, as well as real-time control operations, including controlling the behaviour of smart appliances for demand response. Due to the large number of end-points in distribution systems, real-time monitoring and control in smart grids will require large amounts of data to be managed, which together with the sensitivity of the data gives rise to new data management challenges, including cybersecurity and consumer privacy [2].

Smart meters are expected to be one of the primary sources of real-time monitoring data in smart distribution grids. By measuring and reporting the electricity consumption data of consumers (both industrial and residential) in near real-time, smart meters could enable distribution grid operators to control and optimise the supply and the distribution of electricity, *e.g.,* through real-time distribution system state estimation based on voltages and power flows measured by the smart meters at customers' premises [3]. Furthermore, in the presence of distributed generation and smart appliances, smart meters could enable load balancing through demand response and dynamic protection reconfiguration.

Unfortunately, the data collected by smart meters may also serve for invading consumers' privacy. Several recent works have pointed out that electricity consumption data may allow one to reveal private information, such as household occupancy or economic status [4]–[8]. As a consequence, smart meter data are subject to serious privacy and security concerns[1]. The privacy concerns and their perception within the public have delayed the roll out of smart metering in a number of countries [9], and call for new technical solutions. What makes the problem different from standard data security and privacy is the combination of three factors: the legacy of energy technologies that are based on closed systems, the interweaving with legal and regulatory aspects that introduce additional constraints, and the complex structure of the energy sector, with a variety of interconnected stakeholders, thus requiring more standardised solutions.

In this survey, we present the key privacy issues related to smart meter data collection and management together with its regulatory and policy context (*e.g.,* NIST regulations [10] and EU regulations [11]), we provide an overview of state-of-the-art solutions to preserve privacy, identify shortcomings, provide recommendations and highlight remaining research challenges in order to devise ICT solutions that enable privacy-preserving use of smart meter data, in particular, considering the disparate requirements of the three different uses of the data: billing, operations and value-added services.

A number of recent articles survey security issues in smart grids [12]–[20]. Lu, Wang and Ma [13] provide guidelines on designing security schemes for smart grids. Baumeister [14] reviews and categorises the literature on smart grid security. In [16], Anderson and Fuloria discuss the security economics of electricity metering, while the potential effects of hacking have been reviewed by several security specialists[2]. Mo *et al.* [19] discuss security approaches for smart grids. Privacy challenges in the smart grid have been considered recently in [20]–[26]. Our survey differs from previous surveys on smart grid privacy through (i) considering the privacy aspects of data collection and of data management in an integrated manner while taking into account regulatory aspects [10], (ii) providing an overview of existing solutions, identifying their shortcomings, providing recommendations and highlighting research directions and (iii) contrasting the privacy aspects and

M. R. Asghar is with the Department of Computer Science, The University of Auckland, 1142 Auckland, New Zealand, e-mail: r.asghar@auckland.ac.nz. G. Dán is with the Department of Network and Systems Engineering, School of Electrical Engineering, KTH Royal Institute of Technology, 100 44 Stockholm, Sweden, e-mail: gyuri@ee.kth.se. D. Miorandi is with U-Hopper, 38122 Trento, Italy, e-mail: daniele.miorandi@u-hopper.com I. Chlamtac is with CREATE-NET, 38123 Povo, Trento, Italy, e-mail: chlamtac@create-net.org. This work was partially supported by EIT Digital (http://www.eitdigital.eu/) within the Action Line on Smart Energy Systems. Part of this work appeared in the Proceedings of SmartGridSec 2012 [1].

[1] http://www.smartplanet.com/blog/bulletin/smart-grids-demand-better-protection-from-cyberattacks/

[2] https://www.technologyreview.com/s/420061/hacking-the-smart-grid/

the proposed solutions for the different uses of smart metering data.

The rest of the article is organised as follows. In Section II, we give an overview of smart metering and of the uses of smart meter data. In Section III we describe the legal framework, notions of privacy, and the requirements for privacy-preserving protocols for smart meter data management. In Section IV, we discuss privacy issues related to meter data collection and management, together with their relationship to security. In Section V, we survey privacy-preserving solutions for the three uses of meter data and discuss their shortcomings and potential future research directions. Section VI concludes the article and provides directions for future work.

II. Smart Meter Data and Privacy

The key enablers of smart distribution grids are automated meters and smart meters. We use the term automated meter for a device able to (i) measure consumption of electric energy with a variable time granularity and (ii) report the measured consumption to a Meter Data Management System (MDMS). We use the term smart meter for an automated meter that is additionally able to (iii) receive pricing information or direct load control commands and can (iv) exchange information with smart home appliances, which allows to optimise energy use and to participate in demand response. The data collection and control functionality of automated meters is thus a subset of that of smart meters, and so is the set of privacy issues related to automated meters a subset of the issues related to smart meters.

The metering systems rolled out in most countries record and can transmit measurement data at intervals of about 15 minutes, but hourly and daily reporting are not uncommon [3]. The communication technology used varies depending on the country and population density. Common technologies are Power-Line Communication (PLC), ZigBee [27] and cellular networks. Data are often delivered following a hierarchical model and are processed at, for example, low voltage or medium voltage substations; from there, they are delivered over an IP network to the MDMS. In the future, the data could be delivered from the consumer's premises over 'public' (*i.e.,* non-dedicated) communication networks, such as the Internet.

Figure 1 illustrates the different domains of the smart metering infrastructure, including the customer domain, communications, meter data management, and various services using the metering data.

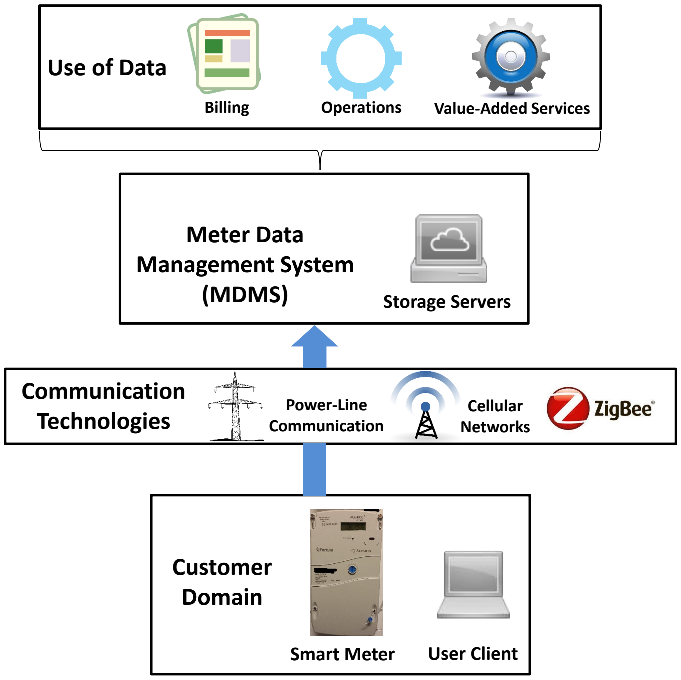

Fig. 1. Different domains of the smart metering infrastructure.

A. Billing, Operations, and Value-Added Services

The measurement data collected by automated and smart meters are one of many sources of information in a smart grid, and are expected to be used in combination with other sources of information collected by the actors that constitute the smart grid, including distribution system operators, the transmission system operator, bulk generation, and electricity markets. Accordingly, the measured data are expected to be used for three kinds of purposes: billing, operations and value-added services. These three kinds of purposes differ significantly in terms of their requirements on metering frequency and accuracy, in terms of the number and locations of consumers whose data are needed, and in terms of the stakeholders.

Billing: The primary use of automated meter data is accurate consumer billing in the presence of dynamic pricing. Since billing typically happens on a monthly basis, the consumption data need not be processed in real-time, but the correctness of billing requires accurate measurement data and accurate time of use information.

Operations: The second use of smart meter data serves for improving the efficiency and the reliability of electricity distribution, especially in the presence of distributed generation. Utilities can use automated meter data for parametrising intelligent Feeder Protection Systems (FPS) at substations in the presence of back-feed due to distributed generation [28], or for improving distribution system State Estimation (SE) and integrated Volt and Var Control (VVC) [29], [30]. They can also use the meter data for detecting faults and for improved automated Fault Location, Isolation and Service Restoration (FLISR). These uses of metering data require real-time or near real-time data processing, but with different granularities and possibly different measurement accuracy. While FPS, SE and

[3] We refer to the "Final Guidelines of Good Practice on Regulatory Aspects of Smart Metering for Electricity and Gas" issued in Feb. 2011 by the European Regulators Group for Electricity and Gas (ERGEG), accessible at: http://www.energy-regulators.eu/portal/page/portal/EER_HOME/EER_PUBLICATIONS/CEER_PAPERS/Customers/Tab2/E10-RMF-29-05_GGP_SM_8-Feb-2011.pdf, and to EANDIS Response to ERGEG's Public Consultation Paper on Draft Guidelines of Good Practice on Regulatory Aspects of Smart Metering for Electricity and Gas at the Council of European Energy Regulators (CEER), accessible at: http://www.ceer.eu/portal/page/portal/EER_HOME/EER_CONSULT/CLOSEDPUBLIC~CONSULTATIONS/CUSTOMERS/Smartmetering/RR/GGPSmartMetering_EANDIS.pdf

VVC may be based on aggregate meter data from a feeder or a section of a feeder, FLISR may need individual meter data for accurate fault location identification, islanding control and restoration planning.

Beyond ancillary and reliability services, information from smart appliances, delivered by smart meters, could be used for improving demand forecasts. For instance, an already active appliance could advertise its future power demand or an appliance programmed to become active at a certain time could advertise its expected power demand curve as a function of dynamic pricing information. Demand forecasts and the controllability of smart appliances can enable the use of demand response for ancillary services, *i.e.,* they can allow an operator to reduce demand by switching off appliances in case of supply scarcity. Such real-time direct demand response requires real-time demand information.

Value-Added Services: Consumers, operators and third party service providers could leverage smart meter data for providing various value-added services, for example, services for the management and for the diagnostics of electric appliances. Value-added energy services could be provided for free or for a fee, depending on the business model of the third party, and could accelerate the transformation of the electricity market [31].

Management services could serve for reducing the energy bill of consumers by giving guidelines (or even control) for economic demand response, *i.e.,* for scheduling the demand as a function of the predicted electricity prices. Controllable loads could include appliances such as washing machines or dishwashers, and home/building energy management systems including Heating, Ventilation and Air Conditioning (HVAC) [32]. As an example, an energy management service provider could provide economic demand response for a fee commensurate to the achieved energy savings. Such economic demand response services could be coordinated with ancillary demand response, and could allow for personalised tariffs depending on customers' engagement in ancillary demand response programs.

Diagnostics and maintenance services could be provided to consumers to identify anomalous consumption patterns, such as appliances consuming excessive energy that should be replaced or appliances close the end of their life cycles. Such services could be implemented by correlating the consumption data of consumers with similar profiles irrespective of their proximity [33]. As an example, an equipment vendor may provide a free service for the lifecycle management of its products in return for statistical information.

To summarise, the requirements of the three uses differ in terms of the frequency at which data need to be measured and collected (real-time vs. batch), in terms of the granularity of the data that are needed (single appliance vs. household vs. group of households) and in terms of the geographical proximity of the consumers whose data are needed.

B. Automated and Smart Meters Collect Personal Data

The data collected by automated meters and by smart meters may serve for a fourth, unintended purpose: to invade consumers' privacy. There has been significant recent work showing that individual appliances (based on their load signatures [34]) can be identified from the detailed analysis of energy consumption traces [35]–[42]. Frequent meter readings can also be used to infer the occupancy of a household, and data mining algorithms can also be used to invade the privacy of consumers in more sophisticated ways, *e.g.,* by revealing their life-styles and economic status [4]–[8]. Recent work also shows that fine enough measurements could reveal consumers' interests as well, *e.g.,* Greveler *et al.* [43] show that they can estimate the displayed TV channel from the electricity usage profile with a sampling period of 0.5 s. Privacy is thus a serious concern, and calls for a data governance framework tailor-made for the smart grid.

III. Privacy and Security Requirements

It should be no surprise in light of the above results that smart meter data has received significant attention both by legislation and by the research community.

A. Privacy Legislation for Smart Meter Data

There is now a consensus that smart meter data should be managed according to the provisions foreseen for "personal data". Recently, the European Data Protection Supervisor issued an opinion on the usage of smart meters' data [4], stating *"Stakeholders must be aware that processing of personal data in the context of smart grids/smart metering will have to fully comply with the national legislation transposing the relevant EU legislation, including Directive 95/46/EC, and – to the extent applicable – the e-Privacy Directive"*[5].

According to the current EU policy[6], the collection of personal data is forbidden unless selectively allowed by law. This includes the case of explicit legitimation, *i.e.,* when the entity collecting personal can demonstrate that the data are necessary for achieving the specific purpose. As an example, a Distribution System Operator (DSO) could demonstrate that smart metering data are necessary for preserving a societal interest (*e.g.,* the stability of the power grid). However, even when allowed, the collection of personal data is subject to limitation of purpose. Personal data collected for one specific purpose (*e.g.,* billing) cannot be used for a different purpose (*e.g.,* profiling); every additional purpose requires a separate legitimation.

In the U.S., there is no federal regulation on the privacy of smart meter data in place, although NIST provided guidelines on privacy aspects in 2010[7]. In California and other states, additional regulations have been established, making the landscape rather fragmented and inhomogeneous[8]. To address this

[4] http://www.edps.europa.eu/EDPSWEB/webdav/site/mySite/shared/Documents/Consultation/Opinions/2012/12-06-08_Smart_metering_EN.pdf

[5] Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications) (OJ L 201, 31.07.2002, p 37), as amended by Directives 2006/24/EC and 2009/136/EC.

[6] EU directive 95/46/EC

[7] http://csrc.nist.gov/publications/nistir/ir7628/nistir-7628_vol2.pdf

[8] http://epic.org/privacy/smartgrid/smartgrid.html

deficiency, a voluntary code of conduct for utilities and for third parties was recently released by the U.S. Department of Energy [44], with guidelines similar to those in EU legislation.

B. Two Notions of Privacy

It is not only the legal framework for privacy that is heterogeneous, but there are also two fundamentally different notions of privacy in the scientific literature.

Cryptographic privacy: Cryptographic privacy requires that the information that is leaked by an algorithm be limited to the information that can be learned from the result of the algorithm [45].

Statistical privacy: Statistical privacy requires that the result of a (possibly randomised) algorithm executed on a data set should not reveal sensitive information about the individuals that constitute the data set, *i.e.,* the objective is to limit the set of inferences that are possible based on the result. A widely used notion in statistical privacy is the notion of differential privacy, which requires that the result of an algorithm be similar when executed on similar data sets, *e.g.,* aggregation should be insensitive to the addition or removal of a consumer in a probabilistic sense [46]. Another notion of privacy is k-anonymity, which requires that for each consumer included in an aggregate, there should be at least $k-1$ other consumers whose data are contained in the aggregate and are indistinguishable [47]. We refer to [48] for an axiomatic discussion of statistical privacy, including differential privacy.

It is important to note that the focus of the two notions of privacy is complementary. To illustrate the difference, consider two queries of electricity consumption data. The first query would request the anonymised hourly electricity consumption data of households over a year. The second query would request blackout locations and durations over the course of the same year. Guaranteeing cryptographic privacy when computing the two queries would ensure the attacker receives anonymised data, but the attacker could link the results of the two queries and infer the consumer locations (and possibly identities) through matching blackout events with reduced household electricity consumption.

Thus, to protect private information, it is important to achieve both notions of privacy. This is even more so if attackers can collude and can manipulate data and protocol information, which leads us to the definition of security requirements and attacker models.

C. Requirements for Privacy-Preserving Protocols for Smart Meter Data Management

Smart meter data management requires a solution that is compliant with privacy regulations; at the same time, it should enable the three kinds of uses of smart meter data discussed in Section II-A. Besides this functional requirement, protocols and solutions for smart meter data management should fulfil a number of security requirements.

- *Confidentiality:* Meter data should not be exposed to unauthorised individuals or processes during transmission (data-in-transit), storage (data-at-rest) and computing (data-in-use). Ensuring confidentiality of data-in-transit, data-at-rest and data-in-use is necessary for achieving cryptographic privacy.
- *Integrity:* The accuracy and correctness of the meter data should be maintained during transmission, storage and computing, and any changes made to the data should be detectable.
- *Authenticity:* The receiver of the meter data should be able to verify the source of the data.
- *Non-Repudiation:* The source of the meter data should not be able to deny that it originated the data. It implies integrity and authenticity.
- *Auditability:* It should be possible to verify whether the response to a request (meter data or computation on meter data) is correct.

How challenging it is to achieve these security requirements while preserving privacy depends significantly on the attacker model. Under the *honest-but-curious* (also called semi-honest) attacker model the adversary is assumed to follow the protocol honestly, *e.g.,* it does not manipulate data. Under the *malicious* attacker model the attacker can deviate from the protocol and can modify protocol messages. A malicious attacker may control multiple meters, or may pretend to have multiple identities, in which case we talk about a Sybil attack.

To address malicious attackers, we thus formulate two additional security requirements that are relevant for solutions that combine data from multiple meters.

- *Non-Malleability:* An attacker should not be able to alter encrypted data without being detected.
- *Sybil Attack Resistance*: The solution should be resilient to meters that present multiple identities.
- *Byzantine Attack Resistance*: The solution should be resilient to colluding meters, *e.g.,* meters that have been compromised.

In what follows, we survey smart meter data management solutions for satisfying these security requirements under the above attacker models.

IV. Privacy Protection for Smart Meter Data under the Trusted Operator Model

The traditional approach to power system data management assumes a private and isolated information infrastructure in which the data are collected and stored by the operator, who is trusted; thus, consumer privacy is ensured. In reality, there is seldom a private and isolated information infrastructure for smart meter data management, and thus even if the operator is trusted, user privacy may be invaded in many different ways. Figure 2 provides an overview of the threats, and Table I shows a summary of the problems, existing solutions and remaining research issues under the trusted operator model, discussed below. The solutions discussed in this section aim at providing cryptographic privacy.

A. Tamper-Resistance of Smart Meters

The most exposed system components, where customer privacy could be invaded, are the meters themselves. Automated and smart meters are typically installed at physically accessible locations, often unprotected, and therefore a smart meter could be the entry point for a number of physical and side channel attacks, made possible by vulnerabilities in the software located in the meter. Typically, a smart meter is considered fully-trusted because it is equipped with a Trusted Platform Module (TPM) [49] for storing cryptographic keys and for performing cryptographic primitives using the keys. If an attacker manages to compromise a smart meter, she could easily get access to the keys and to the measurement data stored in the meter and could invade consumer privacy. Compromising a large number of smart meters by exploiting a common vulnerability of meters connected to the same network infrastructure could enable large-scale real-time privacy invasion, and could serve as an aid for targeted burglary.

Beyond privacy invasion, attacks against individual meters could be motivated by electricity theft, *e.g.,* McLaughlin, Podkuiko and McDaniel describe methods, including password extraction and storage tampering, for adversaries to manipulate consumption data provided by automated meters in [50]. Energy theft using automated meters is not only an academic exercise: automated meters for electricity and for gas were recently found tampered within the U.K., even though meter tampering may result in explosions and even deaths[9].An attacker who can simultaneously compromise many meters by exploiting a common vulnerability could also monetise the compromised meters through providing electricity theft as a service to affected consumers [5]. Alternatively, through switching off a large number of compromised meters or through manipulating the real-time readings from compromised smart meters in a coordinated manner, an attacker could destabilise or could cause operational inefficiency in a distribution system that relies on meter data [51].

Ensuring the integrity of the smart meters is thus of utmost importance for maintaining consumer privacy (and system security). Unfortunately, physically securing smart meters is not an economical solution and, hence, other forms of protection techniques are needed. For avoiding large-scale attacks on smart meters, there is a need for a scalable access control mechanism that would prevent meter compromises but at the same time would allow easy on-site diagnostics and maintenance, akin to using a different password for every meter. For providing security guarantees, there is a need for formal methods for the verification of low-level code typically found in smart meter software, both before the deployment and at firmware upgrades. The verification of such low-level code is rather challenging due to the lack of structure, complicated control flow and a lack of type safety [52]. Moreover, remote attestation and secure logging techniques are needed for enabling large-scale security monitoring (*e.g.,* periodic collection and processing of logs to identify anomalies) and forensic analysis [53]. For a survey of code verification and remote attestation for embedded devices, we refer to [54]–[57].

B. Data Confidentiality and Trust Models

A common approach for preventing unauthorised access to private information is to use encryption. Encryption schemes based on Public Key Infrastructure (PKI) are, for example, widely used for providing confidentiality against third parties during data transmission. In [58], Baumeister investigated what PKI architecture would be most suitable to meet the requirements, in particular scalability and interoperability, of smart grids. His conclusion was that an architecture based on bridge certification authorities would provide a good balance between interoperability, scalability and availability. Using the PKI, a secure communication channel could be established for guaranteeing the confidentiality of the exchanged data if appropriate trust anchors and attribute certificates are put in place and if meters are based on trusted computing platforms, as argued by Metke and Ekl [59]. However, the main issue with the PKI is efficient certificate revocation.

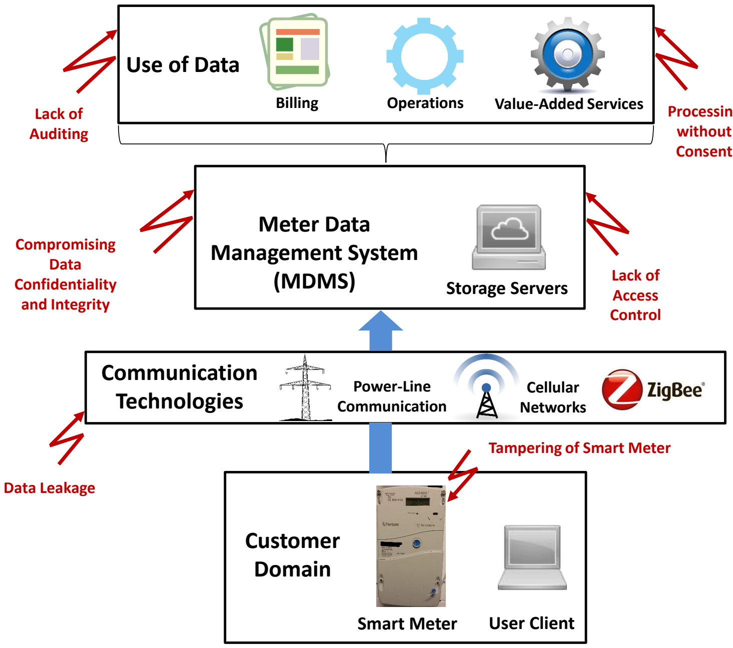

Fig. 2. Overview of threats in the smart metering infrastructure.

State-of-the-art cryptographic solutions, *e.g.,* based on PKI, can keep sensitive data confidential during transmission. Nonetheless, data need also be protected on the storage servers if the servers are located in outsourced environments, such as public or hybrid cloud environments. An outsourced environment is usually considered semi-trusted, assuming that the storage provider is *honest-but-curious* (see, *e.g.,* [60], [61]), *i.e.,* honest to follow the protocols but curious to know about stored and exchanged information.

One approach for providing confidentiality in a semi-trusted environment is to obfuscate the data transmitted to the outsourced storage, *e.g.,* by adding random noise to the data [61], or through an appropriate algebraic transformation of the data [62]. Adding random noise is conceptually simple, may hide critical information from the storage provider, and allows arbitrary computations to be performed on the data. Nonetheless, it may severely affect the accuracy of the computations. In order for perturbation-based obfuscation to be successful, there would be a need of models of the accuracy of computations as a function of the level of confidentiality provided, *i.e.,* the

[9] http://www.bbc.com/news/uk-england-25718447

level of noise. The other approach is based on an algebraic transformation of the input to the optimisation problem and of the result of the computation, *e.g.,* using a randomly chosen linear transformation of the constraint matrix, the objective function coefficients, and a corresponding linear transformation for the solution vector [62]. This approach does not affect the accuracy of the result of the computation, but state-of-the-art techniques only allow solving optimisation problems with linear constraints and linear or quadratic objective functions. Such optimisation problems arise in modern distribution management systems, *e.g.,* for optimal power flow and for distribution state estimation.

An alternative solution is to use encrypted storage. State-of-the-art solutions enable users to perform search over encrypted data at the price of increased computational complexity [60], and so does homomorphic encryption enable computations to be performed over encrypted data at the price of significant computational overhead [63]. While encrypted storage could protect the stored data under the semi-trusted and untrusted storage models, if the units of encrypted smart meter data are small then deterministic encryption schemes [64]–[66] may be easy to break, or could be subject to statistical attacks [67]. A potential solution is to protect the data using probabilistic encryption [60], [68], [69], but this will introduce additional overhead. Unfortunately, the application of existing encryption-based solutions for smart meter data storage is hindered by inefficiency and poor scalability. Furthermore, existing encryption schemes do not allow complex logical expressions to be evaluated over encrypted data, and they assume that there is a single user storing or retrieving the data. Beyond more efficient and scalable algorithms, there is a need for algorithms allowing multi-user access control mechanisms, and for algorithms to ensure resilience for systems that rely on distributed storage servers.

An interesting related issue is whether data access patterns at the MDMS or communication patterns between a meter and the MDMS could reveal private information about a consumer even if the data are encrypted. We are not aware of a reported privacy breach of this kind, but should it be necessary, the data access patterns of algorithms executed on private data can be hidden by Oblivious RAM (ORAM) [70]–[72], while Private Information Retrieval (PIR) [73]–[76] can be used for hiding data access patterns from an untrusted server or database. Finally, steganographic communication [77] can be used for concealing data exchange between two parties.

An outsourced environment could also be considered untrusted. In this case, the integrity of the stored data have to be verified and the results of computations need to be validated [78]. There has been significant recent work on protocols that ensure the integrity of multi-user data stored in untrusted clouds, typically using some form of signed root hash [79]. A straightforward solution for the verification of the results of computations done on non-Byzantine untrusted platforms is to execute the same computation on several different platforms, but this solution is expensive. Solutions for verifiable computation using a single untrusted platform are based on homomorphic encryption [80] and on probabilistically checkable proofs [81], but these solutions are only applicable to computations that can be expressed as a Boolean circuit, and thus they would not be applicable to algorithms used for optimisation.

Although both storage servers may be physically better protected than smart meters, a compromise of the storage server through intrusion or potential malware could lead to privacy violation even under the trusted storage model, as it has happened recently with credit card data[10], [11]. Therefore, the semi-trusted and the untrusted models are not only relevant in the case of outsourcing, but they should be considered as a means of defence in depth for privacy preservation in trusted environments.

C. Consent and Access Control

As discussed in Section III-A, privacy regulation requires that consumers provide explicit consent for each individual use of their private information. Thus, beyond standard uses like billing and operations, any further value-added services, even if provided by the same entity, would require explicit consent from the consumer. To enable a market of third party services, it is therefore essential that consumers have enough information for making informed decisions about regulating access to their private information [82]. Such access control decisions would typically be enforced by the MDMS.

In principle, an authorised entity should be able to access only the requested data, thus following the principle of least privilege [83]. There is a variety of access control mechanisms, including Mandatory Access Control (MAC), Discretionary Access Control (DAC) and Role-Based Access Control (RBAC) [84]. Moreover, flexible access control policies are offered by the eXtensible Access Control Markup Language (XACML) [85], [86]. Nonetheless, flexible access control becomes challenging as we move from coarse-grained to fine-grained access control, because it involves more complex policy specification and incurs significant overhead. At the same time, access control (*e.g.,* [84]) might reveal private information about the sensitive data it aims to protect, in particular if access control policies are enforced in semi-trusted or untrusted environments. A solution discussed in [87] to avoid revealing private information in semi-trusted environments is to use encrypted role-based access control policies; however, it incurs a high computational overhead. The problem of a (usable) specification language to express consent and access control policies that could be used in the context of the smart grid is thus still unsolved.

Even if a specification language existed, giving access to private information for a particular purpose only is technically challenging. Even if a consent is given (as considered in [82]) to a particular use of private information, existing solutions can not verify whether or not the data are processed according to the given consent; thus, it remains an open problem to verify whether or not the data are processed according to the given consent, and to enforce expressive policies in an efficient and scalable manner. Instead of verification, a common approach to ensure privacy is to manipulate the data so that they can only

[10] http://mashable.com/2013/12/19/target-data-breach/

[11] http://mashable.com/2013/07/25/hackers-nasdaq-visa-breach/

be used for the specific purpose for which access is granted. We will discuss such solutions in Section V.

D. Data Integrity and Auditing

Data integrity is essential for accurate metering, and ensuring data integrity may be straightforward using traditional cryptographic solutions, *e.g.,* using PKI-based approaches [58], both for communication and for storage [13]. Nonetheless, protecting data integrity using schemes such as PKI is problematic from a privacy perspective, as it can reveal the identities of consumers to third parties, such as storage providers. For protecting identity information, one may have to rely on anonymous authentication approaches, especially in semi-trusted and untrusted environments [88], [89]. However, anonymous authentication suffers from efficiency and revocation issues. As an alternative, one can consider service-specific solutions, which we review in Section V.

An important problem related to integrity is that of auditing, *i.e.,* verifying whether or not the data received in response to a request are correct. While there exist many solutions for the problem of auditing in general [90]–[94], auditing data without invading privacy is a challenging problem. Recently proposed schemes enable consumers to verify their consumption cost without revealing consumption data to the MDMS [26], [49]. These schemes are based on Pedersen commitments [95] and require that aggregators be honest-but-curious, *i.e.,* they are honest when doing computations (*e.g.,* verifying digital signatures) but are curious to learn about the consumption data. Consequently, these solutions do not support the verification of the computations made by the aggregators. A related problem is that of public auditing of stored data under privacy requirements, *i.e.,* a third party verifying the metering data of a consumer without invading its privacy. A solution for publicly auditing encrypted data on cloud storage was proposed based on a public key-based homomorphic linear authenticator in [96]. Since the solution is based on the verification of blocks of data chosen at random, it is unclear whether such a solution suits smart grids for auditing real-time metering data considering that data are generated continuously, and given the real-time constraints put by the intended use of the data for operations.

V. Service-Specific Privacy Protection under the Non-Trusted Operator Model

The alternative to the trust model considered in Section IV is to consider that operators, utilities and value-added service providers are non-trusted entities. This trust model is mainly motivated by the large number of stakeholders involved in the smart grid ecosystem, whose interests and business models are unknown to consumers. Under this non-trusted operator model, consumer privacy has to be ensured by manipulating the meter data in a way that they can only be used for their intended purposes, for billing, for operations or for one of many value-added services. While we do not discuss it here again, it is important to note that the security of the meters is essential for preserving privacy under this trust model as well. Table II shows a summary of the issues, privacy-preserving solutions and research directions discussed below. Table III provides a detailed comparative analysis of privacy-preserving solutions for smart meter data.

A. Privacy-Preserving Billing

Arguably, the biggest benefit of automated metering is accurate billing in the presence of dynamic pricing. The challenge in privacy-preserving billing is that frequent changes in electricity prices require detailed energy consumption information to enable correct billing, but detailed energy consumption information (*e.g.,* one reading every 15 minutes) may leak private information.

1) Filtering with Energy Storage for Statistical Privacy: One approach to reduce the amount of sensitive information revealed by frequent meter readings is to hide consumption events or the load signatures of individual appliances through charging and discharging an energy storage located at the customers' premises [8], [97]–[101]. The energy storage protects customers' privacy by hiding the use of individual appliances, or by shifting their apparent time of use, and serves in essence as a sort of a low-pass filter of the energy use time series. Nonetheless, the energy storage interacts with dynamic pricing and with demand response [8], and thus it is unclear what would be the optimal battery management strategy and how price and demand predictions can be incorporated in the optimisation.

2) Secure Computation for Cryptographic Privacy: The alternative approach for preserving customers' privacy is to enable energy suppliers to calculate bills without access to individual readings. Jawurek *et al.* [49] propose a scheme based on Pedersen commitments to avoid privacy leakage by introducing a privacy component plugged into the smart meter, which sends only the billing information with the signed commitment to the energy supplier. The commitments are signed by the smart meter and can be verified by the energy supplier. Rial and Danezis [102] extend the idea for calculating bills under a non-linear consumption tariff. Their solution relies on a combination of Groth's integer commitment and the Non-Interactive Zero-Knowledge proof (NIZK) [103], and on the generalised anonymous credential system of Camenish and Lysyanskaya [104]. Motivated by the complexity of the NIZK, they also propose a solution for billing under an affine consumption tariff, which does not rely on the NIZK. The disadvantage of these schemes is that the energy supplier does not get information about the instantaneous power demand, *i.e.,* the information necessary for operations.

A conceptually different technical solution following the same approach is based on a zero-knowledge protocol, proposed by Molina-Markham *et al.* in [7]. In the proposed architecture, the smart meter acts as a prover and the power (or in general energy, such as electricity, gas or water) trace is considered a secret. As the solution is based on a zero-knowledge protocol and relies on homomorphic encryption, it is rather computationally intensive, which may limit its wide scale adoption.

TABLE I
SUMMARY OF PROBLEMS, SOLUTIONS AND ISSUES RELATED TO DATA TRANSMISSION, STORAGE AND PROCESSING UNDER THE TRUSTED OPERATOR MODEL.

Problems	**Solutions**	**Issues**
Tamper-resistance of smart meters	Trusted Platform Module (TPM) [49]	Insider (energy theft) and outsider attacks (exploiting vulnerabilities)
Data confidentiality (trusted environments)	Public Key Infrastructure (PKI) [58], [59]	Efficient certificate revocation
Data confidentiality (semi-trusted environments)	Secure storage [60], encrypted search [69], Oblivious RAM (ORAM) [70]–[72] and Private Information Retrieval (PIR) [73]–[76]	High computational overhead
Data confidentiality (untrusted environments)	Verifiable computation [78]–[81]	Limited applicability
Consumer consent	Automatic consent capturing [82]	No guarantee of data usage per consent
Access control (trusted environments)	Mandatory Access Control (MAC), Discretionary Access Control (DAC), Role-Based Access Control (RBAC) [84] and XACML [85], [86]	Complex policy specification
Access control (semi-trusted environments)	Encrypted RBAC [87]	High computational overhead
Access control (untrusted environments)	None	Open problem
Data integrity (trusted environments)	PKI based approaches [13], [58]	Reveal identities
Data integrity (semi-trusted and untrusted environments)	Anonymous authentication [88], [89]	Efficiency and revocation
Auditing (trusted environments)	Data correctness [90]–[94]	Privacy invasive
Auditing (semi-trusted and untrusted environments)	Homomorphic linear authenticator [96]	Limited applicability

TABLE II
SUMMARY OF SERVICES, PRIVACY-PRESERVING SOLUTIONS AND ISSUES IN THE SMART METERING INFRASTRUCTURE.

Use of Data	**Privacy-Preserving Solutions**	**Issues**
Billing *(Filtering with energy storage)*	Charging and discharging energy storage [8], [97]–[101]	No information about power demand
Billing *(Secure computation)*	Pedersen commitment computed by privacy plug-in [49], Zero-knowledge proof [7], Non-interactive zero-knowledge proof [102]	Computationally intensive
Operations *(with trusted third party)*	Anonymising metering data [105], Spatial and temporal aggregation [106]	Trust in third parties No time of use pricing
Operations *(w/o trusted third party)*	Partially homomorphic encryption [107]–[112], Homomorphic encryption with commitments [26], [113], Homomorphic encryption with DC-Nets [114], DC-Nets [115], [116] Multi-party computation based on wiretap codes [117], Additive noise for differential privacy [106], [116], [118]	Sybil attack, No verifiable computation Or computationally intensive
Value-added services *(for demand response)*	Multi-party computation based on Sharmir's secret sharing [119], Aggregate planned consumption plus additive noise [120]	Limited capabilities

TABLE III

Comparison of privacy-preserving protocols for smart meter data. The table illustrates what uses of data each solution supports including Billing (BL), Operations (OP) and Value-Added Services (VAS). We also show what security properties each solution provides. These security properties are (C)onfidentiality, (I)ntegrity, (AUTH)etication, no(NM)alleability, no(NR)epudiation, (AUD)itability and (ANO)onymity. We indicate if each solution is resistant to (SY)bil attack as well as we provide some remarks.

Solution	**Supports**			**Provides**						**Resistant to**	**Remarks**
	BL	OP	VAS	C	I	AUTH	NM	NR	AUD	SY	
Charging and discharging energy storage [8], [97]–[101]	✓	✗	✗	-	-	-	-	-	-	-	Architectural changes needed
Component for linear tariff [49]	✓	✗	✗	✓	✓	✓	✗	✗	✗	-	Heavy-weight client
Zero-knowledge proof [7]	✓	✓	✗	✓	-	✓	✗	✓	✗	✓	Heavy-weight client/server
Non-interactive zero-knowledge proof [102]	✓	✗	✗	✓	✓	✓	✗	✓	✗	-	Heavy-weight client/server
Anonymising metering data [105]	✓	✓	✗	✓	✓	✓	-	✓	✗	✗	Reliance on neighbours
Spatial and temporal aggregation [106]	-	✓	✗	-	-	-	-	-	-	-	Conceptual model only
Capability-based power management [107]	✗	✓	✗	✓	✗	✗	✗	✗	✗	✓	Requires key sharing
EPPA [108]	✗	✓	✗	✓	✓	✓	✗	✓	✗	✗	Key revocation affects scalability
A decentralised framework for data aggregation [109]	✗	✓	✗	✓	-	-	✗	✓	✗	✗	Key revocation affects scalability
Information aggregation scheme by Li, Luo and Liu [110], [112]	✗	✓	✗	✓	✗	✗	✗	✗	✗	✗	Requires key sharing
A privacy-friendly smart metering architecture [111]	✓	✓	✗	✓	✓	✓	✗	✓	✗	✗	Requires key sharing
Homomorphic encryption with commitments [26], [113]	✓	✓	✗	✓	✓	✓	✗	✓	✓	✗	Requires key sharing
Homomorphic encryption with DC-Nets [114]	✓	✓	✗	✓	✗	✗	✗	✗	✓	✓	Requires key sharing
Symmetric DC-Nets [115], [116]	✗	✓	✗	✓	-	-	-	-	✓	✓	Requires key sharing
Asymmetric DC-Nets [26]	✓	✓	✗	✓	✓	✓	✓	✓	✓	✓	Complex cryptographic primitives
Multi-party computation based on wiretap codes [117]	✗	✓	✗	✓	-	-	-	-	-	✓	Tunable overhead
Additive noise for differential privacy [106], [116], [118]	✗	✓	?	✓	-	-	-	-	-	✓	Noise affects operations & VAS
Multi-party computation based on Sharmir's secret sharing [119]	✗	✓	✗	✓	-	-	-	-	-	✗	Requires all to be involved
Aggregate planned consumption plus additive noise [120]	✗	✓	✗	✓	-	-	-	-	-	✗	Noise-based theoretical model

Although there are a number of proposals for privacy-preserving billing based on a non-linear consumption tariff, apart from the solutions based on an on-site battery, there is yet no solution that would allow providing information to the operator about the instantaneous power demand for supporting operations, *e.g.,* voltage control. Battery-based solutions, at the same time, make certain value-added services infeasible, *e.g.,* the monitoring of anomalous consumption patterns.

B. Privacy-Preserving Operations

The second, important use of smart meter data is improving operational efficiency and safety in distribution grids. Unlike in the case of billing, individual smart meter data may not be necessary for improving operations. Instead, it may be sufficient for a utility to know the instantaneous aggregate power demand and the instantaneous aggregate power supply within areas of the power network. The size of the area and thus the level of aggregation depends on the application, *e.g.,* neighbourhood, substation, district or an entire city, but, of course, the number of consumers aggregated influences the level of privacy protection, and it may also affect the operational efficiency, *e.g.,* in the case of distribution state estimation [29], [30] or system identification [28].

1) Aggregation Algorithms for Cryptographic Privacy: There is a wealth of literature on aggregation algorithms for cryptographic privacy, with and without a trusted third party. The focus of most of these works is on privacy-preserving aggregation, little attention has been paid to the trade-off between privacy-preservation and the usefulness of the data for improving operational efficiency.

a) With a Trusted Third Party: Several solutions for cryptographically privacy-preserving aggregation of meter data rely on a trusted third party [105], [106]. Efthymiou and Kalogridis [105] describe a scheme for anonymising metering data assuming a trusted escrow service to aggregate the smart meter data to be anonymised. In their solution, each smart meter has a non-anonymous client data profile with the utility for billing, and an anonymous data profile with the escrow service that aggregates data from several meters. Bohli *et al.* [106] propose a solution that aggregates the consumption information of consumers in a neighbourhood and sends it to the energy supplier. At the end of each billing period, the solution also submits the aggregate consumption per consumer, but it does not allow time-of-use pricing. Like [105], the solution merely transfers the trust to the neighbourhood gateways. A notable weakness of the approaches relying on a trusted third party is the possibility of the trusted third party being compromised, thus there needs to be a solution that can allow the customers or the operators to identify in real-time any compromise of the trusted third party.

b) Without a Trusted Third Party: Solutions without a trusted third party rely on secure multi-party computation [121], implemented either using some form of homomorphic encryption or using channel codes designed for the wire-tap channel. Homomorphic encryption allows non-trusted third parties to perform operations over encrypted data. Ideally, the third party could apply any number and combination of arithmetic operations (*e.g.,* addition, subtraction, multiplication and division) on the encrypted data, without deteriorating the accuracy of the results. Achieving this is, however, difficult. State-of-the-art homomorphic encryption algorithms fall within one of two categories: partially homomorphic cryptosystems and fully homomorphic cryptosystems. A partially homomorphic cryptosystem exclusively offers either addition or multiplication. The Paillier [122], the Goldwasser-Micali [123] and the Benaloh [124] cryptosystems (an extension of the Goldwasser-Micali cryptosystem) are examples of partially homomorphic cryptosystems that support addition. ElGamal [125] and RSA [126] are examples of partially homomorphic cryptosystems that support multiplication. Unlike partially homomorphic cryptosystems, fully homomorphic cryptosystems support both addition and multiplication. The most recent and probably most complete such cryptosystem is [63], but these schemes are not practical yet due to the high computational overhead incurred by the underlying cryptographic operations.

The use of additive partially homomorphic cryptosystems for aggregation without a trusted third party has been explored in centralised, decentralised and distributed architectures. In EPPA [108] consumption data are aggregated at network gateways, using the Paillier cryptosystem. Ruj and Nayak [109] introduce a solution based on the Paillier cryptosystem combined with attribute-based encryption, in which data are aggregated in a hierarchical manner by network elements, such as home area gateways, building area gateways and remote terminal units. Li *et al.* [110] propose a scheme in which smart meters connected by a mesh network form a tree, and aggregation is done by the smart meters themselves using the Paillier cryptosystem. Common between these works is that meters are assumed to be *honest-but-curious* and malleability is thus not a concern. Instead of aggregating data in the network, Vetter *et al.* [111] propose an architecture in which data are encrypted using a partial homomorphic cryptosystem that not only allows aggregation of encrypted values but also the aggregation of encryption keys. The encrypted data are stored in a database grouped by region, and the energy provider can query and decrypt aggregate data using the aggregated encryption key. However, this scheme is not secure against known-plaintext attacks.

Solutions based on symmetric and asymmetric DC-Nets were considered in [26], [115]. DC-Nets were introduced by Chaum [127] for computing the Boolean or secret values. They rely on a temporary secret shared among participants, and provide unconditional security, but are sensitive to disruption attacks, *i.e.,* a malicious attacker can render the result of the computation useless. Aggregation protocols based on DC-Nets were introduced and analysed in [115], including a low overhead aggregation protocol that establishes shared secrets using public keys.

Contrary to Chaum's symmetric DC-Nets, in asymmetric DC-Nets [128] participants use their permanent private keys for encryption and the aggregator uses the sum of the private keys, which it is assumed to know, for decrypting the aggregate value. Asymmetric DC-Nets do not provide unconditional security, or perfect forward secrecy, but they allow the aggregator to verify the aggregated value and individual values (at the price of sacrificing privacy). [26] considers solutions based on symmetric and asymmetric DC-Nets, and argues that asymmetric DC-Nets are a generalisation of the solutions based on additive partial homomorphic cryptosystems. Yet, asymmetric DC-Nets require more complex cryptographic primitives (including exponentiation and multiplication [26]) as compared to symmetric DC-Nets, which require XOR or addition [127].

An example of a secure multi-party computation scheme that does not rely on homomorphic encryption is the solution based on wiretap codes presented in [117], which allows to compute linear functions of data distributed in a network of smart meters and has an overhead that grows linearly in the number of meters.

2) Providing Statistical Privacy: An important issue in the case of solutions based on multi-party computation is their potential vulnerability to Sybil attacks, *i.e.,* to colluding meters whose aim is to reveal private data of other meters. A number of privacy-preserving aggregation schemes that are resilient against colluding meters have been formulated and evaluated using the notion of differential privacy, which has found widespread use in solutions for privacy-preserving aggregation without a trusted third party [106], [116], [118], through adding random noise to the measurement data.

Under the assumption that the noise is normally distributed, the authors in [106] concluded that achieving the desired level of differential privacy would require too large aggregation groups for the solution to be practical. A more widely used distribution for the noise is the Laplace distribution, which was used in [116] to devise a protocol that relies on data exchange between the meters. The protocol is robust to faulty nodes, but malicious meters may be able to make the data irrecoverable.

A critical aspect of achieving differential privacy by adding random noise is the potential impact of the noise on smart grid applications, such as state estimation, restoration, dynamic relay configuration or VVC. A first step in this direction is recent work that quantified the trade-off between differential privacy and state estimation accuracy under Gaussian and Laplacian noise on a single feeder [118]. Nonetheless, it is unclear if these results can be generalised for topologies that are more general, and for a wider range of applications, including optimal power flow, VVC and FLISR.

Another critical aspect of achieving differential privacy by adding random noise is whether algorithms exist that allow revealing a power consumption time series in real time without allowing an attacker to leverage the temporal correlation between subsequent samples for invading privacy. An approach based on filtering and adaptive sampling was proposed in [129] for epidemic and for traffic data, but it is unclear if such an approach would preserve consumer privacy if applied to smart meter data. Finally, it is unclear whether it is possible to add random noise in a way that it does not affect the correctness of billing.

3) Privacy Economics: An alternative approach to aggregation would be to provide an economic incentive to customers for sharing frequent meter reading data with the operator. Such a market-based solution would allow each customer to decide about the reporting frequency individually, depending on the time-of-day, the sensitivity of its activities, and the financial incentive, thus setting a price on privacy [130]. At the same time, the operator could adjust the economic incentive to the value of the data received from the customers in improving the efficiency and safety of its operation. We are not aware of a framework exploring this interesting direction.

C. Value-Added Services

Privacy-preserving value-added services have received much less attention than operations and billing. The value-added service that has received the most attention is privacy-preserving demand response [119], [120], [131]. The scheme presented in [131] assumes a trusted entity to which customers submit their bids in the form of the power demand they would be willing to shed and the corresponding price. The solution presented in [119] uses secure multi-party computation based on Shamir's secret sharing scheme and relies on a set of schedulers that can be honest-but-curious. As an alternative, an iterative scheme that assumes that customers exchange aggregate consumption plans with additive noise was proposed in [120].

Besides demand-response, value-added services could aim at identifying appliance level anomalies, or could optimise the electricity consumption of a household, and are related to Non-Intrusive Load Monitoring (NILM). Unlike NILM, whose primary objective is load disaggregation [35], value-added services would aim at providing value to the consumer based on characteristics of a household's energy consumption, and would presumably require metering data with different resolutions. It is worthwhile to point out that the feasibility of such value-added services does not contradict privacy-preserving operations, as services may not need complete time series, *e.g.,* high frequency data may be sufficient for identifying faulty rectifiers [132]. An important aspect of outlier detection is that a characterisation of normal behaviour may not be available. Outliers may thus have to be identified using unsupervised machine learning, *e.g.,* privacy-preserving outlier detection using distance-based methods was shown to be possible using distributed algorithms [133].

The naive solution to enable such value-added services would be to transmit several down-sampled versions of the same data, and protect each version with a key. A more sophisticated solution is to use a hierarchical representation, *e.g.,* by recursively applying a wavelet transform on the low pass sub-band [134]. These representations need a hierarchical key management scheme, which ensures that the customer can provide a key to a value-added service provider that gives access to the right representations [135]. An alternative solution would be to use a source-coding paradigm similar to Multiple Description Coding (MDC) used for loss resilient audio and video coding [136]. An MDC-like scheme would create several representations of equal importance, and a service provider could use an arbitrary k-subset of these representations to obtain an encoding of the data with sufficient accuracy for performing the service.

Research in this area would benefit from the definition of the resolution that various value-added services would need and whether data are needed continuously or only occasionally. For example, a high-pass filtered version of a household's power consumption may be sufficient for services that identify malfunctioning equipment, but would not reveal the household's average consumption. If high-resolution data are needed occasionally, it may be more efficient to use a separate, public communication infrastructure for data exchange, as in this way the metering infrastructure would require significantly less bandwidth, and privacy-preserving technologies developed for operation and billing need not be extended to high frequency data for value-added services. Such an engineering solution would, however, increase the cost of the data collection infrastructure, and raises security concerns due to the direct connection of smart meters to public communication infrastructures.

Another alternative would be to implement value-added services on the customers' private computing platforms, such as their mobile phones. Doing so would allow private data to be kept locally, but even assuming that the data can be delivered to the devices in a privacy-preserving manner, there are still a number of potential issues that need to be dealt with. First, this solution would require value-added service providers to deploy their algorithms on customer-owned devices, running the risk of their algorithms to be stolen. Second, many value-added services would likely employ some form of machine learning and would thus be computationally intensive. Third, it is likely that many value-added services would rely on comparing data from different customers, in which case distributed privacy preserving algorithms would be needed for implementing value-added services. Addressing these issues will require progress in the area of energy-efficient and privacy-preserving distributed machine learning algorithms.

VI. Conclusion and Future Directions

In this article, we have surveyed the state-of-the-art on smart meter data privacy. Focusing on the three uses of smart meter data, and its privacy aspects, we have reviewed cryptographic solutions for ensuring privacy-preserving management of smart meter data under the trusted operator model, and privacy-preserving solutions for data processing under the non-trusted operator model. Despite the wealth of solutions proposed in the literature, there are several open problems in smart meter data privacy. In the following, we highlight some of these exciting open problems.

Meter Data Management. Considering business incentives and the regulatory framework, it is very likely that smart meter data will be managed and processed at utility-managed or at third party data centres. Secure and privacy-preserving management of smart meter data will therefore be fundamental, both under the trusted environment model and under the untrusted environment model. Considering the trusted storage model, an interesting avenue for research would be the investigation of the fundamental limits of obfuscation-based solutions for confidentiality-preserving computing. The alternative to obfuscation-based solutions would be scalable encrypted storage that allows complex logical expressions to be evaluated, possibly on streaming data, with multi-user access control and with support for resilience. Results in this area, especially encrypted storage that allows machine learning algorithms to be executed on private data, would allow a variety of value-added services, but could also find applications in other domains. A third exciting avenue for research under the trusted operator model is the problem of access control and consent management, *i.e.,* flexible privacy-preserving fine-grained access control and the related verification of whether or not the data are processed according to the consent given.

Considering the untrusted storage model, the most fundamental problem is that of verifiable computation for general optimisation problems. Since smart meter data are generated and will be used in real-time for operations, solutions should have a low complexity and may possibly have to support distributed execution. Finally, due to the financial and operational safety implications of data manipulation, privacy-preserving public auditing of real-time data will be a basic requirement with no known available solution.

Privacy-Preserving Billing, Operations and Value-Added Services. The solutions presented in Section V (also listed in Tables II and III) focus on solving privacy issues for one or two uses of smart meter data, typically billing and/or operations. While many service-specific solutions address integrity and confidentiality, they do not ensure auditability, non-repudiation and resistance against Sybil attacks, and the fundamental limitations of statistical privacy for smart meter data are not very well understood either. It is worthwhile to note that auditability, non-repudiation and resistance against Sybil attacks are particularly challenging to achieve if consumer privacy is to be preserved. Whether there exists a solution that supports all uses of data and satisfies all security requirements is still an open question.

Furthermore, there has been surprisingly little work on the definition of value-added services beyond economic demand-response, and consequently very little attention has been paid to privacy-preserving solutions. A promising approach to privacy-preserving data management for value-added services could be to use solutions developed for audio and video source coding, but further research is needed to understand whether such an approach could as well support billing and operations using a single data management infrastructure. An alternative, but similarly exciting direction is to investigate the possibility of value-added services implemented on customers' premises, which requires advances both in the area of privacy-preserving distributed machine learning algorithms and in the area of code protection.

Economic Models of Privacy. Economic models of privacy have been developed for a variety of contexts, but we are not aware of works in the area of privacy economics for smart meter data. Privacy economics for digital economies is particularly interesting due to information asymmetry caused by that consumers are not well informed about what data are used and for what purpose. Work in the area of privacy economics could develop mechanisms for information sharing that would allow consumers to make rational decisions, and could develop models of how people could be compensated for revealing their private information that would then allow maximising social welfare. Addressing this issue would require both utility theoretic models of privacy [137] and game theoretic models of consumer-operator interaction.

Acknowledgement

György Dán was partly funded by the EU H2020 SUCCESS project, grant agreement No. 700416, and by the MSB CERCES project.

References

[1] M. R. Asghar and D. Miorandi, “A holistic view of security and privacy issues in smart grids,” in *Smart Grid Security*. Springer, 2013, pp. 58–71.

[2] Executive Office of the President of the US, “A policy framework for the 21st century grid: Enabling our secure energy future,” http://www.whitehouse.gov/sites/default/files/microsites/ostp/nstc-smart-grid-june2011.pdf, National Science and Technology Council, June 2011, last accessed: June 11, 2016.

[3] X. Han, S. You, F. Thordarson, D. V. Tackie, S. M. stberg, O. M. Pedersen, H. W. Bindner, and N. C. Nordentoft, “Real-time measurements and their effects on state estimation of distribution power system,” in *Proc. of IEEE PES Innovative Smart Grid Technologies (ISGT) Europe*, Oct. 2013, pp. 1–5.

[4] G. Wood and M. Newborough, “Dynamic energy-consumption indicators for domestic appliances: environment, behaviour and design,” *Energy and Buildings*, vol. 35, no. 8, pp. 821–841, 2003.

[5] P. McDaniel and S. McLaughlin, “Security and privacy challenges in the smart grid,” *IEEE Security and Privacy Mag.*, vol. 7, no. 3, pp. 75–77, 2009.

[6] E. Quinn, “Privacy and the new energy infrastructure,” *Available at SSRN http://ssrn.com/abstract=1370731*, 2009.

[7] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, “Private memoirs of a smart meter,” in *Proc. of ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building (BuildSys)*, 2010, pp. 61–66.

[8] G. Kalogridis, R. Cepeda, S. Denic, T. Lewis, and C. Efthymiou, “Elecprivacy: Evaluating the privacy protection of electricity management algorithms,” *IEEE Trans. on Smart Grid*, vol. 2, no. 4, pp. 750–758, Dec 2011.

[9] C. Cuijpers and B.-J. Koops, “Smart metering and privacy in Europe: Lessons from the Dutch case,” in *European Data Protection: Coming of Age*, S. Gutwirth, R. Leenes, P. de Hert, and Y. Poullet, Eds. Springer, 2013, pp. 269–293.

[10] A. Lee and T. Brewer, “Smart grid cyber security: Strategy and requirements,” https://www.smartgrid.gov/sites/default/files/doc/files/NISTIR_7628_Draft_1_Smart_Grid_Cyber_Security_Strategy_Requi_200902.pdf, NIST, Sep. 2009, last accessed: June 11, 2016.

[11] Enisa, “Smart grid security,” http://www.enisa.europa.eu/activities/Resilience-and-CIIP/critical-infrastructure-and-services/smart-grids-and-smart-metering/smart-grid-security-related-initiatives, ENISA, March 2012, last accessed: June 11, 2016.

[12] N. Komninos, E. Philippou, and A. Pitsillides, “Survey in smart grid and smart home security: Issues, challenges and countermeasures,” *IEEE Communications Surveys Tutorials*, vol. 16, no. 4, pp. 1933–1954, 2014.

[13] X. Lu, W. Wang, and J. Ma, “Authentication and integrity in the smart grid: An empirical study in substation automation systems,” *International Journal of Distributed Sensor Networks*, Apr. 2012.

[14] T. Baumeister, “Literature review on smart grid cyber security,” https://csdl-techreports.googlecode.com/svn/trunk/techreports/2010/10-11/10-11.pdf, Collaborative Software Development Laboratory, Department of Information and Computer Sciences, University of Hawaii, Tech. Rep., December 2010, last accessed: January 29, 2016.

[15] H. Khurana, M. Hadley, N. Lu, and D. A. Frincke, “Smart-grid security issues,” *IEEE Security and Privacy Mag.*, vol. 8, pp. 81–85, 2010.

[16] R. Anderson and S. Fuloria, “On the security economics of electricity metering,” in *Proc. of Workshop on Economics of Internet Security (WEIS)*, 2010.

[17] W. Wang and Z. Lu, “Cyber security in the smart grid: Survey and challenges,” *Computer Networks*, vol. 57, no. 5, pp. 1344 – 1371, 2013.

[18] Z. Fan, P. Kulkarni, S. Gormus, C. Efthymiou, G. Kalogridis, M. Sooriyabandara, Z. Zhu, S. Lambotharan, and W. H. Chin, “Smart grid communications: Overview of research challenges, solutions, and standardization activities,” *IEEE Communications Surveys Tutorials*, vol. 15, no. 1, pp. 21–38, 2013.

[19] Y. Mo, T.-H. Kim, K. Brancik, D. Dickinson, H. Lee, A. Perrig, and B. Sinopoli, “Cyber-physical security of a smart grid infrastructure,” *Proc. of the IEEE*, vol. 100, no. 1, pp. 195–209, Jan. 2012.

[20] K. Sharma and L. M. Saini, “Performance analysis of smart metering for smart grid: An overview,” *Renewable and Sustainable Energy Reviews*, vol. 49, pp. 720 – 735, 2015.

[21] S. Finster and I. Baumgart, “Privacy-aware smart metering: A survey,” *IEEE Communications Surveys Tutorials*, vol. 17, no. 2, pp. 1088–1101, 2015.

[22] J. Liu, Y. Xiao, S. Li, W. Liang, and C. L. P. Chen, “Cyber security and privacy issues in smart grids,” *IEEE Communications Surveys Tutorials*, vol. 14, no. 4, pp. 981–997, 2012.

[23] A. Cavoukian, J. Polonetsky, and C. Wolf, “SmartPrivacy for the smart grid: embedding privacy into the design of electricity conservation,” *Identity in the Information Society*, vol. 3, no. 2, pp. 275–294, 2010.

[24] S. Finster and I. Baumgart, “Privacy-aware smart metering: A survey,” *IEEE Communications Surveys Tutorials*, vol. 16, no. 3, pp. 1732–1745, 2014.

[25] Z. Yang, S. Yu, W. Lou, and C. Liu, “p^2 : Privacy-preserving communication and precise reward architecture for v2g networks in smart grid,” *IEEE Trans. on Smart Grid*, vol. 2, no. 4, pp. 697–706, Dec. 2011.

[26] F. Borges, *On privacy-preserving protocols for smart metering systems*. Springer, 2017.

[27] D.-M. Han and J.-H. Lim, “Design and implementation of smart home energy management systems based on ZigBee,” *IEEE Trans. on Consumer Electronics*, vol. 56, no. 3, pp. 1417 –1425, Aug. 2010.

[28] J. Peppanen, M. J. Reno, M. Thakkar, S. Grijalva, and R. G. Harley, “Leveraging ami data for distribution system model calibration and situational awareness,” *IEEE Trans. on Smart Grid*, vol. 6, 2015.

[29] K. Samarakoon, J. Wu, J. Ekanayake, and N. Jenkins, “Use of delayed smart meter measurements for distribution state estimation,” in *Proc. of IEEE PES General Meeting*, 2011.

[30] P. A. Pegoraro, J. Tang, J. Liu, F. Ponci, A. Monti, and C. Muscas, “Pmu and smart metering deployment for state estimation in active distribution grids,” in *Proc. of IEEE ENERGYCON*, 2012.

[31] C. E. Kontokosta, “Energy disclosure, market behavior, and the building data ecosystem,” *Annals of the New York Academy of Sciences*, vol. 1295, no. 1, pp. 34–43.

[32] K. Vatanparvar, Q. Chaun, and M. Al Faruque, “Home energy management as a service over networking platforms,” in *Proc. of IEEE PES ISGT*, 2015.

[33] S. Karnouskos, “Smart houses in the smart grid and the search for value-added services in the cloud of things era,” in *Proc. of IEEE Intl. Conf. on Industrial Technology*, 2013.

[34] H. Y. Lam, G. S. K. Fung, and W. K. Lee, “A novel method to construct taxonomy of electrical appliances based on load signatures,” *IEEE Trans. on Consumer Electronics*, vol. 53, no. 2, pp. 653–660, May 2007.

[35] N. Batra, J. Kelly, O. Parson, H. Dutta, W. Knottenbelt, A. Rogers, A. Singh, and M. Srivastava, “NILMTK: An Open Source Toolkit for Non-intrusive Load Monitoring,” in *Proc. of ACM Intl. Conf. on Future Energy Systems (e-Energy)*, 2014.

[36] K. Anderson, A. Ocneanu, D. Benitez, D. Carlson, A. Rowe, and M. Berges, “BLUED: a fully labeled public dataset for Event-Based Non-Intrusive load monitoring research,” in *Proc. of ACM KDD Workshop on Data Mining Applications in Sustainability (SustKDD)*, Beijing, China, Aug. 2012.

[37] M. Zeifman, “Disaggregation of home energy display data using probabilistic approach,” *IEEE Trans. on Consumer Electronics*, vol. 58, no. 1, pp. 23–31, Feb. 2012.

[38] J. Z. Kolter and M. J. Johnson, “Redd: A public data set for energy disaggregation research,” in *Proc. of ACM KDD Workshop on Data Mining Applications in Sustainability (SustKDD), San Diego, CA*, Aug. 2011.

[39] J. Liang, S. K. K. Ng, G. Kendall, and J. W. M. Cheng, “Load signature study-part ii: Disaggregation framework, simulation, and applications,” *IEEE Trans. on Power Delivery*, vol. 25, no. 2, pp. 561–569, Apr. 2010.

[40] M. Baranski and J. Voss, “Genetic algorithm for pattern detection in NIALM systems,” in *IEEE Intl. Conf. on Systems, Man and Cybernetics*, vol. 4, Oct. 2004, pp. 3462–3468.

[41] A. Prudenzi, “A neuron nets based procedure for identifying domestic appliances pattern-of-use from energy recordings at meter panel,” in *IEEE PES Winter Meeting*, vol. 2, 2002, pp. 941–946.

[42] G. Hart, “Nonintrusive appliance load monitoring,” *Proc. of the IEEE*, vol. 80, no. 12, pp. 1870–1891, Dec. 1992.

[43] U. Greveler, B. Justus, and D. Loehr, "Multimedia content identification through smart meter power usage profiles," *Computers, Privacy and Data Protection*, 2012.

[44] U.S. Department of Energy, "Data privacy and the smart grid: A voluntary code of conduct," 2015.

[45] R. Küsters, E. Scapin, T. Truderung, and J. Graf, *Extending and Applying a Framework for the Cryptographic Verification of Java Programs*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 220–239.

[46] A. Friedman and A. Schuster, "Data mining with differential privacy," in *Proc. of ACM Intl. Conf. on Knowledge Discovery and Data Mining (KDD)*, 2010, pp. 493–502.

[47] L. Sweeney, "k-anonymity: A model for protecting privacy," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 5, pp. 557–570, 2002.

[48] B.-R. Lin and D. Kifer, "Information measures in statistical privacy and data processing applications," *ACM Trans. Knowl. Discov. Data*, vol. 9, no. 4, pp. 28:1–28:29, Jun 2015.

[49] M. Jawurek, M. Johns, and F. Kerschbaum, *Plug-In Privacy for Smart Metering Billing*. Springer, 2011, pp. 192–210.

[50] S. McLaughlin, D. Podkuiko, and P. McDaniel, "Energy theft in the advanced metering infrastructure," in *Critical Information Infrastructures Security (CRITIS)*, E. Rome and R. Bloomfield, Eds. Springer, 2010, vol. 6027, pp. 176–187.

[51] A. Teixeira, G. Dán, H. Sandberg, R. Berthier, R. B. Bobba, and A. Valdes, "Security of smart distribution grids: Data integrity attacks on integrated Volt/VAR control and countermeasures," in *Proc. of American Control Conference (ACC)*, Jun. 2014.

[52] M. Balliu, M. Dam, and R. Guanciale, "Automating information flow analysis of low level code," in *Proc. of ACM Conf. on Computer and Communications Security (CCS)*, Nov. 2014.

[53] C. G. M. LeMay, "Cumulative attestation kernels for embedded systems," in *Proc. of European Symposium on Research in Computer Security (ESORICS)*, Sep. 2009.

[54] A. Seshadri, A. Perrig, L. van Doorn, and P. Khosla, "Swatt: software-based attestation for embedded devices," in *Proc. of IEEE Symp. on Security and Privacy*, May 2004, pp. 272–282.

[55] A. Seshadri, M. Luk, A. Perrig, L. van Doorn, and P. Khosla, "Scuba: Secure code update by attestation in sensor networks," in *Proc. of ACM Workshop on Wireless Security (WiSe)*, 2006, pp. 85–94.

[56] C. Castelluccia, A. Francillon, D. Perito, and C. Soriente, "On the difficulty of software-based attestation of embedded devices," in *Proc. of ACM Conf. on Computer and Communications Security (CCS)*, 2009, pp. 400–409.

[57] D. Perito and G. Tsudik, *Proc. of European Symposium on Research in Computer Security (ESORICS)*. Springer, 2010, ch. Secure Code Update for Embedded Devices via Proofs of Secure Erasure, pp. 643–662.

[58] T. Baumeister, "Adapting PKI for the smart grid," in *Proc. of IEEE SmartGridComm*, October 2011, pp. 249–254.

[59] A. Metke and R. Ekl, "Security technology for smart grid networks," *IEEE Trans. on Smart Grid*, vol. 1, no. 1, pp. 99–107, Jun. 2010.

[60] M. R. Asghar, G. Russello, B. Crispo, and M. Ion, "Supporting complex queries and access policies for multi-user encrypted databases," in *Proc. of ACM Cloud Computing Security Workshop (CCSW)*, 2013, pp. 77–88.

[61] O. Vukovic, G. Dán, and R. B. Bobba, "Confidentiality-preserving obfuscation for cloud-based power system contingency analysis," in *Proc. of IEEE SmartGridComm*, Nov. 2013.

[62] A. R. Borden, D. K. Molzahn, P. Ramanathan, and B. C. Lesieutre, "Confidentiality-preserving optimal power flow for cloud computing," in *Allerton Control Conference*, 2012.

[63] C. Gentry, "A fully homomorphic encryption scheme," Ph.D. dissertation, Stanford University, 2009.

[64] R. A. Popa, C. M. S. Redfield, N. Zeldovich, and H. Balakrishnan, "CryptDB: Processing queries on an encrypted database," *Commun. of the ACM*, vol. 55, no. 9, pp. 103–111, Sep. 2012.

[65] A. Boldyreva, N. Chenette, Y. Lee, and A. ONeill, "Order-preserving symmetric encryption," in *Proc. of Advances in Cryptology (EUROCRYPT)*, ser. LNCS, A. Joux, Ed. Springer, 2009, vol. 5479, pp. 224–241.

[66] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu, "Order preserving encryption for numeric data," in *Proc. of ACM Intl. Conf. on Management of Data (SIGMOD)*, 2004, pp. 563–574.

[67] A. Ceselli, E. Damiani, S. De Capitani Di Vimercati, S. Jajodia, S. Paraboschi, and P. Samarati, "Modeling and assessing inference exposure in encrypted databases," *ACM Trans. Inf. Syst. Secur.*, vol. 8, no. 1, pp. 119–152, Feb. 2005.

[68] C. Dong, G. Russello, and N. Dulay, "Shared and searchable encrypted data for untrusted servers," *Journal of Computer Security*, vol. 19, no. 3, pp. 367–397, 2011.

[69] D. Boneh, G. Di Crescenzo, R. Ostrovsky, and G. Persiano, "Public key encryption with keyword search," in *Proc. of Advances in Cryptology (EUROCRYPT)*, C. Cachin and J. Camenisch, Eds. Springer, 2004, vol. 3027, pp. 506–522.

[70] E. Stefanov, M. van Dijk, E. Shi, C. Fletcher, L. Ren, X. Yu, and S. Devadas, "Path ORAM: An extremely simple oblivious ram protocol," in *Proc. of ACM Conf. on Computer and Communications Security (CCS)*, 2013, pp. 299–310.

[71] O. Goldreich and R. Ostrovsky, "Software protection and simulation on oblivious RAMs," *Journal of the ACM*, vol. 43, no. 3, pp. 431–473, May 1996.

[72] R. Ostrovsky, "Efficient computation on oblivious RAMs," in *Proc. of ACM Symposium on Theory of Computing (STOC)*, 1990, pp. 514–523.

[73] N. Borisov, G. Danezis, and I. Goldberg, "DP5: A Private Presence Service," in *Proc. of Privacy Enhancing Technologies (PETS)*, Jun. 2015.

[74] S. Yekhanin, "Private Information Retrieval," *Commun. of the ACM*, vol. 53, no. 4, pp. 68–73, Apr. 2010.

[75] P. Williams and R. Sion, "Usable PIR," in *Proc. of Network and Distributed System Security Symp. (NDSS)*, Feb. 2008.

[76] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan, "Private information retrieval," *Journal of the ACM*, vol. 45, no. 6, pp. 965–981, Nov. 1998.

[77] N. Johnson and S. Jajodia, "Exploring steganography: Seeing the unseen," *IEEE Computer*, vol. 31, no. 2, pp. 26–34, Feb. 1998.

[78] G. Dán, R. B. Bobba, G. Gross, and R. H. Campbell, "Cloud computing for the power grid: From service composition to assured clouds," in *Proc. of Usenix HotClouds*, Jun. 2013.

[79] A. Shraer, C. Cachin, A. Cidon, I. Keidar, Y. Michalevsky, and D. Shaket, "Venus: Verification for untrusted cloud storage," in *Proc. of ACM Cloud Computing Security Workshop (CCSW)*, October 2010.

[80] R. Gennaro, C. Gentry, and B. Parno, "Non-interactive verifiable computing: Outsourcing computation to untrusted workers," in *Proc. of Advances in Cryptology (CRYPTO)*, 2010, pp. 465–482.

[81] S. Setty, R. McPherson, A. Blumberg, and M. Walfish, "Making argument systems for outsourced computation practical (sometimes)," in *Proc. of Network and Distributed System Security Symposium (NDSS)*, 2012.

[82] M. R. Asghar and G. Russello, "ACTORS: A goal-driven approach for capturing and managing consent in e-health systems," in *IEEE Intl. Symp. on Policies for Distributed Systems and Networks (POLICY)*, Jul. 2012, pp. 61–69.

[83] J. Saltzer and M. Schroeder, "The protection of information in computer systems," *Proc. of the IEEE*, vol. 63, no. 9, pp. 1278–1308, 1975.

[84] R. S. Sandhu, E. J. Coyne, H. L. Feinstein, and C. E. Youman, "Role-based access control models," *IEEE Computer*, vol. 29, no. 2, pp. 38–47, 1996.

[85] S. Godik, A. Anderson, B. Parducci, P. Humenn, and S. Vajjhala, "OASIS eXtensible Access Control 2 Markup Language (XACML) 3," Tech. rep., OASIS, Tech. Rep., 2002.

[86] R. Yavatkar, D. Pendarakis, and R. Guerin, "IETF RFC 2753: A framework for policy based admission control," Jan. 2000, available at: http://docstore.mik.ua/rfc/rfc2753.html.

[87] M. R. Asghar, "Privacy preserving enforcement of sensitive policies in outsourced and distributed environments," Ph.D. dissertation, University of Trento, Italy, Dec. 2013, http://eprints-phd.biblio.unitn.it/1124/.

[88] J. Camenisch, S. Hohenberger, M. Kohlweiss, A. Lysyanskaya, and M. Meyerovich, "How to win the clonewars: Efficient periodic n-times anonymous authentication," in *Proc. of ACM Conf. on Computer and Communications Security (CCS)*, 2006, pp. 201–210.

[89] P. P. Tsang, M. H. Au, A. Kapadia, and S. W. Smith, "PEREA: Towards practical ttp-free revocation in anonymous authentication," in *Proc. of ACM Conf. on Computer and Communications Security (CCS)*, 2008, pp. 333–344.

[90] P. T. Devanbu, M. Gertz, C. U. Martel, and S. G. Stubblebine, "Authentic third-party data publication," in *Proc. of IFIP TC11/WG11.3 Conf. on Database Security: Data and Application Security, Development and Directions*, 2001, pp. 101–112.

[91] H. Pang, A. Jain, K. Ramamritham, and K.-L. Tan, "Verifying completeness of relational query results in data publishing," in *Proc. of ACM Intl. Conf. on Management of Data (SIGMOD)*, 2005, pp. 407–418.

[92] B. Parno, J. Howell, C. Gentry, and M. Raykova, "Pinocchio: Nearly practical verifiable computation," in *Proc. of IEEE Symp. on Security and Privacy*, 2013, pp. 238–252.

[93] M. Backes, D. Fiore, and R. M. Reischuk, "Verifiable delegation of computation on outsourced data," in *Proc. of ACM Conf. on Computer and Communications Security (CCS)*, 2013, pp. 863–874.

[94] S. Benabbas, R. Gennaro, and Y. Vahlis, "Verifiable delegation of computation over large datasets," in *Proc. of Conf. on Advances in Cryptology (CRYPTO)*, 2011, pp. 111–131.

[95] T. P. Pedersen, *Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing*. Springer, 1992, pp. 129–140.

[96] C. Wang, S. S. Chow, Q. Wang, K. Ren, and W. Lou, "Privacy-preserving public auditing for secure cloud storage," *IEEE Trans. on Computers*, vol. 62, no. 2, pp. 362–375, 2013.

[97] M. Backes and S. Meiser, "Differentially private smart metering with battery recharging," in *Data Privacy Management and Autonomous Spontaneous Security*, J. Garcia-Alfaro, G. Lioudakis, N. Cuppens-Boulahia, S. Foley, and W. M. Fitzgerald, Eds. Springer, 2014, pp. 194–212.

[98] O. Tan, D. Gunduz, and H. Poor, "Increasing smart meter privacy through energy harvesting and storage devices," *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 31, no. 7, pp. 1331–1341, Jul. 2013.

[99] S. McLaughlin, P. McDaniel, and W. Aiello, "Protecting consumer privacy from electric load monitoring," in *Proc. of ACM Conf. on Computer and Communications Security (CCS)*, 2011, pp. 87–98.

[100] D. Varodayan and A. Khisti, "Smart meter privacy using a rechargeable battery: Minimizing the rate of information leakage," in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*, May 2011, pp. 1932–1935.

[101] G. Kalogridis, C. Efthymiou, S. Denic, T. Lewis, and R. Cepeda, "Privacy for smart meters: Towards undetectable appliance load signatures," in *Proc. of IEEE SmartGridComm*, Oct. 2010, pp. 232–237.

[102] A. Rial and G. Danezis, "Privacy-preserving smart metering," in *Proc. of ACM Workshop on Privacy in the Electronic Society (WPES)*, 2011, pp. 49–60.

[103] J. Groth, "Non-interactive zero-knowledge arguments for voting," in *Proc. of Intl. Conf. on Applied Cryptography and Network Security (ACNS)*. Springer, 2005, pp. 467–482.

[104] J. Camenisch and A. Lysyanskaya, "A signature scheme with efficient protocols," in *Proc. of Intl. Conf. on Security in Communication Networks (SCN)*. Springer, 2002, pp. 268–289.

[105] C. Efthymiou and G. Kalogridis, "Smart grid privacy via anonymization of smart metering data," in *Proc. of IEEE SmartGridComm*, October 2010, pp. 238–243.

[106] J. Bohli, C. Sorge, and O. Ugus, "A privacy model for smart metering," in *Proc. of IEEE Intl. Conf. on Communications Workshops (ICCW)*, 2010, pp. 1–5.

[107] D. Seo, H. Lee, and A. Perrig, "Secure and efficient capability-based power management in the smart grid," in *Parallel and Distributed Processing with Applications Workshops (ISPAW), 2011 Ninth IEEE International Symposium on*, May 2011, pp. 119–126.

[108] R. Lu, X. Liang, X. Li, X. Lin, and X. Shen, "EPPA: An efficient and privacy-preserving aggregation scheme for secure smart grid communications," *IEEE Trans. on Parallel and Distributed Systems*, vol. 23, no. 9, pp. 1621–1631, Sept. 2012.

[109] S. Ruj and A. Nayak, "A decentralized security framework for data aggregation and access control in smart grids," *IEEE Transactions on Smart Grid*, vol. 4, no. 1, pp. 196–205, March 2013.

[110] F. Li, B. Luo, and P. Liu, "Secure and privacy-preserving information aggregation for smart grids," *Int. Journal of Security and Networks*, vol. 6, no. 1, pp. 28–39, 2011.

[111] B. Vetter, O. Ugus, D. Westhoff, and C. Sorge, "Homomorphic primitives for a privacy-friendly smart metering architecture," in *Proc. of International Conference on Security and Cryptography (SECRYPT)*, 2012, pp. 102–112.

[112] F. Li, B. Luo, and P. Liu, "Secure information aggregation for smart grids using homomorphic encryption," in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, Oct 2010, pp. 327–332.

[113] F. Borges, D. Demirel, L. Bock, J. A. Buchmann, and M. Mühlhäuser, "A privacy-enhancing protocol that provides in-network data aggregation and verifiable smart meter billing," in *2014 IEEE Symposium on Computers and Communications (ISCC)*, June 2014, pp. 1–6.

[114] Z. Erkin and G. Tsudik, *Private Computation of Spatial and Temporal Power Consumption with Smart Meters*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 561–577.

[115] K. Kursawe, G. Danezis, and M. Kohlweiss, *Privacy-Friendly Aggregation for the Smart-Grid*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 175–191.

[116] G. Ács and C. Castelluccia, "I have a dream! (differentially private smart metering)," *Information Hiding*, pp. 118–132, 2011.

[117] R. Thobaben, G. Dán, and H. Sandberg, "Wiretap codes for secure multi-party computation," in *Proc. of IEEE GlobeCom Workshop on Trusted Communications with Physical Layer Security (TCPLS)*, 2014.

[118] H. Sandberg, G. Dán, and R. Thobaben, "Differentially private state estimation in distribution networks with smart meters," in *Proc. of IEEE Conf. on Decision and Control (CDC)*, Dec. 2015.

[119] C. Rottondi and G. Verticale, "Privacy-friendly appliance load scheduling in smart grids," in *Proc. of IEEE SmartGridComm*, Oct. 2013, pp. 420–425.

[120] C. E. Rottondi, A. Barbato, and G. Verticale, "A privacy-friendly game-theoretic distributed scheduling system for domestic appliances," in *Proc. of IEEE SmartGridComm*, Nov. 2014.

[121] K. B. Frikken, "Algorithms and theory of computation handbook," M. J. Atallah and M. Blanton, Eds. Chapman & Hall/CRC, 2010, ch. Secure Multiparty Computation, pp. 14–14. [Online]. Available: http://dl.acm.org/citation.cfm?id=1882723.1882737

[122] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *Proc. of Advances in Cryptology (EUROCRYPT)*, J. Stern, Ed. Springer, 1999, pp. 223–238.

[123] S. Goldwasser and S. Micali, "Probabilistic encryption," *Journal of Computer and System Sciences*, vol. 28, no. 2, pp. 270 – 299, 1984.

[124] J. Benaloh and D. Tuinstra, "Receipt-free secret-ballot elections (extended abstract)," in *Proc. of ACM Symposium on Theory of Computing (STOC)*, 1994, pp. 544–553.

[125] T. ElGamal, "A public key cryptosystem and a signature scheme based on discrete logarithms," in *Proc. of Advances in Cryptology (CRYPTO)*, ser. LNCS, G. Blakley and D. Chaum, Eds. Springer, 1985, pp. 10–18.

[126] R. L. Rivest, A. Shamir, and L. Adleman, "A method for obtaining digital signatures and public-key cryptosystems," *Commun. of the ACM*, vol. 21, no. 2, pp. 120–126, Feb. 1978.

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

[128] F. Borges, J. Buchmann, and M. Mhlhuser, "Introducing asymmetric DC-Nets," in *Communications and Network Security (CNS), 2014 IEEE Conference on*, Oct 2014, pp. 508–509.

[129] L. Fan and L. Xiong, "Real-time aggregate monitoring with differential privacy," in *Proc. of ACM Intl. Conf. on Information and Knowledge Management (CIKM)*, 2012, pp. 2169–2173.

[130] A. Acquisti, C. R. Taylor, and L. Wagman, "The economics of privacy," *Journal of Economic Literature*, vol. 52, no. 2, 2016.

[131] A. J. Paverd, A. Martin, and I. Brown, "Privacy-enhanced bi-directional communication in the smart grid using trusted computing," in *Proc. of IEEE SmartGridComm*, Nov. 2014.

[132] R. Isermann, *Fault-Diagnosis Applications: Model-Based Condition Monitoring*. Springer-Verlag, 2011.

[133] J. Vaidya and C. Clifton, "Privacy-preserving outlier detection," in *Proc. of IEEE ICDM*, 2004.

[134] D. Engel, "Wavelet-based load profile representation for smart meter privacy," in *Proc. of IEEE PES Innovative Smart Grid Technologies (ISGT)*, Feb. 2013, pp. 1–6.

[135] D. Peer, C. Engel and S. Wicker, "Hierarchical key management for multi-resolution load data representation," in *Proc. of IEEE SmartGridComm*, Nov. 2014.

[136] V. Goyal, "Multiple description coding: compression meets the network," *IEEE Signal Processing Mag.*, vol. 18, no. 5, pp. 74–93, Sep. 2001.

[137] A. Acquisti, L. John, and G. Loewenstein, "What is privacy worth?" *The Journal of Legal Studies*, vol. 42, no. 2, pp. 249–274, June 2013.

Vitae

Muhammad Rizwan Asghar is a Senior Lecturer in the Department of Computer Science at The University of Auckland in New Zealand. Previously, he was a Post-Doctoral Researcher at international research institutes including the Center for IT-Security, Privacy, and Accountability (CISPA) at Saarland University in Germany and CREATE-NET in Trento Italy. He received his Ph.D. degree from the University of Trento, Italy in 2013. As part of his Ph.D. programme, he was a Visiting Fellow at the Stanford Research Institute (SRI), California, USA. He obtained his M.Sc. degree in Information Security Technology from the Eindhoven University of Technology (TU/e), The Netherlands in 2009. His research interests include access control, applied cryptography, security, privacy, cloud computing and distributed systems.

György Dán is an Associate Professor at KTH Royal Institute of Technology, Stockholm, Sweden. He received the M.Sc. in Computer Engineering from the Budapest University of Technology and Economics, Hungary in 1999, the M.Sc. in Business Administration from the Corvinus University of Budapest, Hungary in 2003, and the Ph.D. in Telecommunications from KTH in 2006. He worked as a Consultant in the field of access networks, streaming media and videoconferencing from 1999 to 2001. He was a Visiting Researcher at the Swedish Institute of Computer Science in 2008, a Fulbright research scholar at University of Illinois at Urbana-Champaign in 2012-2013, and an invited Professor at EPFL in 2014-2015. He was co-chair of the Cyber Security and Privacy Symposium at IEEE SmartGridComm 2014, and is an Area Editor of Elsevier Computer Communications. His research interests include the design and analysis of content management and computing systems, game theoretical models of networked systems, and cyber-physical system security in power systems.

Daniele Miorandi is Executive VP for R&D at U-Hopper and Chief Research Officer at Thinkinside. He received a Ph.D. in Communications Engineering from University of Padova, Italy, in 2005. His current research interests include modelling and performance analysis of large-scale networked systems, ICT platforms for socio-technical systems and distributed optimisation for smart grids. Dr. Miorandi has co-authored more than 130 papers in internationally refereed journals and conferences. He serves on the Steering Committee of various international events (WiOpt, Autonomics, ValueTools), for some of which he was a co-founder (Autonomics and ValueTools). He also serves on the TPC of leading conferences in the networking and computing fields. He is a member of ACM, ISOC and EAI. This work was carried out when he was lead scientist at CREATE-NET.

Imrich Chlamtac is the President of CREATE-NET, the Bruno Kessler Professor at the University of Trento, Italy and the President of the European Alliance for Innovation (EAI). Further, he has held various honorary and chaired professorships in USA and Europe including the Distinguished Chair in Telecommunications Professorship at the University of Texas at Dallas, Sackler Professorship at Tel Aviv University and University Professorship at the Technical University of Budapest. In the past, he was with Technion and UMass, Amherst, DEC Research. Dr. Imrich Chlamtac has made significant contribution to various networking technologies as scientist, educator and entrepreneur. Dr. Chlamtac is the recipient of multiple awards and recognitions including Fellow of the IEEE, Fellow of the ACM, Fulbright Scholar, the ACM Award for Outstanding Contributions to Research on Mobility and the IEEE Award for Outstanding Technical Contributions to Wireless Personal Communications. Dr. Chlamtac published close to four hundred refereed journal, book, and conference articles and is listed among ISIs Highly Cited Researchers in Computer Science. Dr. Chlamtac is the co-author of four books, inluding the first book on Local Area Networks (1980) and the Amazon.com best seller and IEEE Editor's Choice Wireless and Mobile Network Architectures, published by John Wiley and Sons (2000). Dr. Chlamtac has widely contributed to the scientific community as founder and Chair of ACM Sigmobile, founder and steering committee chair of some of the lead conferences in networking, including ACM Mobicom, IEEE/SPIE/ACM OptiComm, CreateNet Mobiquitous, CreateNet WiOpt, IEEE/CreateNet Broadnet, IEEE/CreateNet Tridentcom and IEEE/CreateNet Securecomm conferences. Dr. Chlamtac also serves as the founding Editor in Chief of the ACM/URSI/Springer Wireless Networks (WINET), the ACM/Springer Journal on Special Topics in Mobile Networks and Applications (MONET).