Privacy-Enhanced Bi-Directional Communication in the Smart Grid using Trusted Computing

Andrew Paverd
Department of Computer Science
University of Oxford
andrew.paverd@cs.ox.ac.uk

Andrew Martin
Department of Computer Science
University of Oxford
andrew.martin@cs.ox.ac.uk

Ian Brown
Oxford Internet Institute
University of Oxford
ian.brown@oii.ox.ac.uk

***Abstract*—Although privacy concerns in smart metering have been widely studied, relatively little attention has been given to privacy in bi-directional communication between consumers and service providers. Full bi-directional communication is necessary for incentive-based demand response (DR) protocols, such as demand bidding, in which consumers bid to reduce their energy consumption. However, this can reveal private information about consumers. Existing proposals for privacy-enhancing protocols do not support bi-directional communication. To address this challenge, we present a privacy-enhancing communication architecture that incorporates all three major information flows (network monitoring, billing and bi-directional DR) using a combination of spatial and temporal aggregation and differential privacy. The key element of our architecture is the Trustworthy Remote Entity (TRE), a node that is singularly trusted by mutually distrusting entities. The TRE differs from a trusted third party in that it uses Trusted Computing approaches and techniques to provide a technical foundation for its trustworthiness. A automated formal analysis of our communication architecture shows that it achieves its security and privacy objectives with respect to a previously-defined adversary model. This is therefore the first application of privacy-enhancing techniques to bi-directional smart grid communication between mutually distrusting agents.**

I. Introduction

It is widely acknowledged that there are privacy concerns associated with smart energy meters. Various privacy-enhancing protocols and systems have been proposed to mitigate the risk of private information being inferred from frequent energy consumption measurements. Experience has shown that consumers do not always trust service providers such as the energy supplier or distribution network operator (DNO) with these fine-grained consumption measurements [1]. Even if service providers follow the defined protocols, they might still be perceived as honest-but-curious (HBC) adversaries attempting to learn private information about consumers [2]. In addition to these privacy concerns, there are numerous security threats that must also be taken into account in smart grid communication protocols. Most privacy-enhancing protocols that have been proposed focus on two main information flows: monitoring and billing. In the monitoring flow, consumers send frequent consumption measurements to the DNO to allow fine-grained monitoring of the distribution network. In the billing flow, these frequent measurements are sent to the energy supplier to facilitate price-based demand response schemes such as dynamic pricing. In dynamic pricing, the price of energy varies with time to encourage consumers to reduce consumption during periods of high demand. The communication of energy price information is not included in these information flows as it is usually sent via a broadcast channel. Both monitoring and billing are therefore uni-directional information flows.

Demand response (DR) is defined as: *"Changes in electric usage by end-use customers from their normal consumption patterns in response to changes in the price of electricity over time, or to incentive payments designed to induce lower electricity use at times of high wholesale market prices or when system reliability is jeopardized"* [3]. As shown in this definition and by Albadi et al. [4], there are two types of DR approaches: price-based (e.g. dynamic pricing) and incentive-based DR. An example of incentive-based DR is *demand bidding* [4] in which consumers interact with the Demand Side Manager (DSM) as follows: When a shortage in supply is expected, the DSM notifies all consumers. Consumers send *bids* to the DSM stating the amount of consumption they are willing to reduce and the desired incentive price for this reduction. The DSM selects the winning bids and communicates its decision to the individual consumers. Afterwards the respective incentives are credited to successful bidders. Demand bidding therefore requires full bi-directional communication between consumers and the DSM and this constitutes a third primary information flow.

Compared to dynamic pricing, the main advantage of demand bidding is that it provides a closed feedback loop allowing the DSM to estimate the level of DR participation for the next time period. With dynamic pricing there is no guarantee that raising the price will reduce demand and by the time the effects are reported by smart meters, the demand situation might have worsened. In contrast, demand bidding allows the level of DR participation to be determined and optimized. Standards such as OASIS Energy Interoperation (Ei) [5] specify data models for demand bidding. Although initially targeted at industrial consumers, demand bidding can also be applied to residential consumers. In residential settings, a home energy management system or feature-rich smart meter would place bids and control appliances according to a user-defined policy.

However, it has been shown that these bids can be used to infer private information about consumers [2][6][7]. The magnitude and timing of a particular bid could reveal the use of a particular type of system (e.g. charging a plug-in electric vehicle after arriving home). If multiple bids can be linked to a specific consumer, these can be used to infer behavioural patterns. Any deviation from these patterns could also reveal private information [2]. The challenge is that existing proposals for privacy-enhanced smart metering do not support bi-directional communication.

To address this challenge, we present a unified communication architecture incorporating all three primary information flows. The key element of our architecture is the Trustworthy Remote Entity (TRE), a communication node that is singularly trusted by mutually distrusting entities. The TRE is an intermediary in the communication path between consumers and service providers. The TRE enhances consumers' privacy using a combination of spatial and temporal aggregation techniques and facilitates privacy-preserving bi-directional communication for demand bidding protocols.

Although the TRE performs a similar role to a trusted third party, the two concepts are fundamentally different. By definition, a trusted third party is trusted by the relying parties without proof. In contrast, the TRE provides a technical mechanism through which its trustworthiness can be verified. We present a mechanism for establishing trust in the TRE using Trusted Computing (TC) approaches and technologies. Due to the unique characteristics of the smart grid, we propose that existing TC components and approaches, such as the Trusted Platform Module (TPM) and remote attestation, can be leveraged to provide meaningful security guarantees.

Previous work has shown how some proposed smart grid protocols have failed to meet their security and privacy requirements [8]. To avoid this, we have formally analysed the security and privacy properties of our protocols using an enhanced version of the Casper/FDR protocol analysis tool [9]. Our key contributions are:

- A privacy-enhancing communication architecture incorporating all three smart grid information flows, utilizing Trustworthy Remote Entities (TREs).
- An approach for establishing trust in the TRE using tools and approaches from the field of Trusted Computing.
- An automated formal analysis showing that our communication architecture improves upon existing protocols, particularly with respect to bi-directional communication.

II. Related Work

A. Privacy in Smart Metering

For network monitoring, measurements do not need to be attributable to individual consumers. It is sufficient for the DNO to receive aggregated consumption data from a group of smart meters. Some proposals involve anonymizing or pseudonymizing individual measurements [10][11][12][13]. However, it has been shown that in some cases these can be de-pseudonymized [8]. Anonymization is not directly suitable for demand bidding because the incentives cannot be credited to anonymous bidders. Other proposals use *spatial aggregation* in which measurements from a group of consumers are added together to hide each individual's contribution. Proposed spatial aggregation mechanisms include homomorphic encryption [14][15], data perturbation [16][17] and secret sharing [18]. However, these cannot be used for demand bidding because the DSM is unable to select and notify individual bidders.

For billing purposes, measurements must be attributable to individual, named consumers. Some privacy-preserving billing protocols use *temporal aggregation* in which measurements from a single consumer are aggregated over time. The energy supplier has a significant financial interest in ensuring that this aggregation has been performed honestly. It has been proposed that this can be achieved using verifiable computation techniques [19][20]. However, temporal aggregation cannot be used for demand bidding because the bidding interactions must take place in real time. Danezis et al. [21] have proposed an improved temporal aggregation approach that uses data perturbation based on differential privacy [22] to enhance consumers' privacy whilst still providing some real time feedback to the supplier. Although the data is attributable to named consumers and approximate individual measurements are available in real time, this approach is still not ideal for demand bidding because perturbation of bids could result in unacceptably large overall errors. Furthermore, the DSM requires accurate data in order to select successful bidders and credit them with the relevant incentives. Some proposals for monitoring or billing have included trusted third parties [16][23][24] but none of these have addressed the challenge of bi-directional communication.

B. Privacy in Demand Response

Early research efforts have begun to investigate privacy-enhancing techniques for DR applications. Rottondi and Verticale [7] have proposed the use of Multi-Party Computation (MPC) to facilitate privacy-friendly appliance load-scheduling. Although it addresses a similar problem, their architecture is designed for collaborative scheduling rather than incentive-based DR. The most similar work is that of Karwe and Strüker [6] who investigated a demand bidding protocol using a different threat scenario. They showed that an untrusted intermediary between the consumers and the DSM could compromise consumers' privacy and they proposed a mitigation strategy [6]. We focus on the complementary threat scenario in which the third party is provably trustworthy whilst all other agents are mutually distrusting. We have previously analysed the privacy issues in bi-directional DR communication using different types of adversary models and suggested that these issues could be mitigated using TREs [2]. This paper is the fulfilment of that suggestion and, to the best of our knowledge, the first work to address privacy concerns in bi-directional DR communication between mutually distrusting agents.

III. Communication Architecture

We first describe the baseline system model and the current security and privacy threats. We then present our privacy-enhancing communication architecture in terms of the three main information flows and discuss alternative approaches and implementation considerations.

A. Baseline System Model

In the set of all consumers $\mathcal{C}$, each consumer $c \in \mathcal{C}$ has a feature-rich smart meter or home energy management system capable of bi-directional communication. At time $t \in \mathbb{N}$, c produces a consumption measurement m_t^c and sends this to the DNO (monitoring) and to the supplier (billing). For dynamic pricing, the supplier periodically broadcasts the prevailing price per unit p_t to all consumers but this is not included in the billing information flow because it is a broadcast message. When incentive-based DR is required, the DSM notifies consumers and invites bids. At time t, each consumer c may generate a DR bid $(bid\text{-}q_t^c, bid\text{-}p_t^c)$ consisting of at least a bid quantity and bid price per unit and send this to the DSM. Depending on the specific implementation, the bid quantity could be expressed as energy or power and additional information such as the starting time and duration could be included with the bid. However, this additional information is not always present (e.g. for fixed duration hour-ahead bids) so is omitted for the remainder of this paper. Once the bids have been received, the DSM replies to individual consumers indicating acceptance of their bids. We refer to this bi-directional communication between the consumer and DSM as the DR information flow. In all cases, the communicating entities identify and authenticate themselves to one another. Although real implementations could involve more complicated protocols, the above model captures the fundamental elements of the three main information flows.

In this baseline model, there are various threats to consumers' privacy as well as the overall security of the system. We focus on the threat model defined in [2]. In this model, it is assumed that a limited number of consumers are adversarial and will submit false measurements (a type of false data injection attack). It is also assumed that all service providers could be honest-but-curious (HBC) adversaries who will follow the defined protocol but will attempt to learn private information about consumers from any received messages [2].

B. Enhanced System Model

In our privacy-enhancing communication architecture architecture, all communication between consumers and service providers passes through a TRE. For each information flow, the TRE performs specific information processing tasks as described in the following subsections. In all cases, communication with the TRE takes place over secure authenticated channels providing confidentiality and integrity protection with respect to external adversaries as well as strong mutual authentication. This could be achieved using Transport Layer Security (TLS) with mutual authentication. Although we describe the functionality of a single TRE, we envisage that there will be a network of TREs distributed throughout the grid, each providing identical functionality.

C. Network Monitoring

In the network monitoring information flow, the TRE performs spatial aggregation over a group of consumers and applies data perturbation to the result to achieve differential privacy [22]. Consumers are divided into aggregation groups $g \subset \mathcal{C}$ where $\mathcal{G}$ is the set of all groups on a particular TRE. The aggregation groups are dynamically defined by the DNO such that each group $g \in \mathcal{G}$ represents a sector within the distribution network.

Every 15 or 30 minutes, at time t, consumers send individual measurements m_t^c to the TRE. The TRE first performs bounds checking to mitigate against false data injection attacks. Measurements that exceed a consumer's installed capacity will be excluded from the aggregation and an alert will be raised. For each aggregation group, the TRE computes the sum of the measurements and adds random noise according to the Laplace distribution, Lap(λ), which has the density function $h(y) \propto exp(-|y|/\lambda)$ (mean $= 0$, standard deviation $= \lambda$) [25]. The result is sent to the DNO:

$$\text{TRE} \rightarrow \text{DNO:} \quad \left(\sum_{c \in g} m_t^c\right) + Y \; ; \quad \text{where: } Y \sim \text{Lap}(1/\epsilon)$$

This mechanism is therefore ϵ-indistinguishable [25]. The addition of random noise necessary to mitigate against a variant of the *set-difference* attack [26] in which the DNO creates two overlapping aggregation groups that differ by a single consumer in order to learn that individual's consumption. The *sensitivity* of the added noise is calibrated to mask the presence of absence of any single consumer in the aggregate [25]. All consumers in a particular group must connect to the same TRE. Consumers' privacy is technically preserved if $|g| \geq 2$ but in practice, larger aggregation groups would be used. As $|g|$ increases, the percentage error introduced by the random noise decreases. In all practical implementations, this error will be less than other errors such as those caused by electrical losses in the distribution network. The maximum $|g|$ depends on implementation details such as the bandwidth and computational capacity of the TRE. This approach achieves the same outcome as other spatial aggregation techniques [14][15][16][17][18] without requiring any modification to the smart meters and only minimal configuration changes at the DNO. Specifically, this approach does not increase the number of messages sent by the consumers.

D. Billing

In the billing information flow, temporal aggregation is used to preserve the level of privacy available before smart meters. At time t the supplier notifies the TRE of the current energy price p_t which the TRE then broadcasts to consumers. By verifying that p_t was sent by the TRE, consumers are assured that this is the price that will be applied. Consumers send measurements m_t^c to the TRE which performs bounds checking and adds them to the consumer's running total:

$$bill_t^c \; = \; bill_{t-1}^c + (m_t^c \times p_t)$$

At the end of the billing period ($t = t\text{-}end$), the TRE sends each consumer's aggregated total to the energy supplier and resets the running total:

$$\text{TRE} \rightarrow \text{Supplier:} \quad bill_{t\text{-}end}^c \; ; \quad bill_{t\text{-}end}^c = 0$$

The temporal aggregation period is dynamically defined by the supplier but must exceed the minimum value specified by the regulator and enforced by the TRE to protect privacy (e.g. weeks or months). It is not necessary to apply differential privacy in this case because the supplier cannot define overlapping time periods and thus cannot learn anything other than the temporal aggregate. The maximum temporal aggregation period is again implementation-dependent. This achieves the same result as other privacy-preserving billing methods [21][19][20] without requiring modifications to the smart meters or increasing the number of messages sent by consumers. The TRE can combine the temporal aggregation for billing purposes with the spatial aggregation for monitoring since both use the same individual measurements as inputs.

E. Demand Response

Due to the requirement for full bi-directional communication in the DR information flow, techniques such as spatial or temporal aggregation cannot be used directly. For example, in demand bidding, the bids cannot be spatially aggregated over multiple consumers because each bid contains both quantity and price information. Furthermore, residential consumers may only have the ability to reduce demand by a specific amount (e.g. disconnecting a particular load) so bids cannot be partially accepted. In our architecture, the TRE combines the functionality of a privacy proxy with temporal aggregation techniques as shown in Figure 1. When the DSM creates a new DR event, the TRE notifies the consumers and participating consumers submit bids to the TRE:

$$\forall c \in \mathcal{C}: \quad c \rightarrow \text{TRE:} \quad (bid\text{-}q_t^c, bid\text{-}p_t^c)$$

To mitigate against false bid injection, consumers must authenticate themselves to the TRE and the TRE performs bounds checking on all bids. The TRE sends the DSM a set of *pseudo-bids* corresponding to the consumers' bids:

$$\forall c \in \mathcal{C}: \quad \text{TRE} \rightarrow \text{DSM:} \quad (pseudo\text{-}q_t^c, pseudo\text{-}p_t^c)$$

Each pseudo-bid includes a single-use anonymous identifier that can only be linked to the original bid by the TRE. This differs from pseudonymization in which the same pseudonym would be used for all bids from a particular consumer, thus allowing linkability between bids. From the DSM's perspective, the TRE appears to be a large aggregated load that submits multiple bids for each DR event. The DSM can therefore use its existing processes and algorithms to select a set of accepted specific pseudo-bids, $\mathcal{A}$ and notify the TRE which in turn notifies the individual consumers. If this information flow were viewed in isolation, the TRE would still not enable full demand bidding because the incentives could not be credited to individual consumers. However, since our architecture incorporates all three information flows, the TRE can credit the consumers' internal aggregated bills:

$$\forall\ c \in \mathcal{A} \quad : \quad bill_t^c \ = \ bill_{t-1}^c - (bid\text{-}q_t^c \times bid\text{-}p_t^c)$$

These incentives are therefore included in the temporal aggregation of the billing data thus preventing them from being used to link bids to individual consumers. If required, the TRE could also verify that successful bidders have complied with their bid obligations based on their consumption measurements. This protocol ensures that the DSM is unable to link bids to individual consumers and is therefore unable to detect if specific consumers have placed bids.

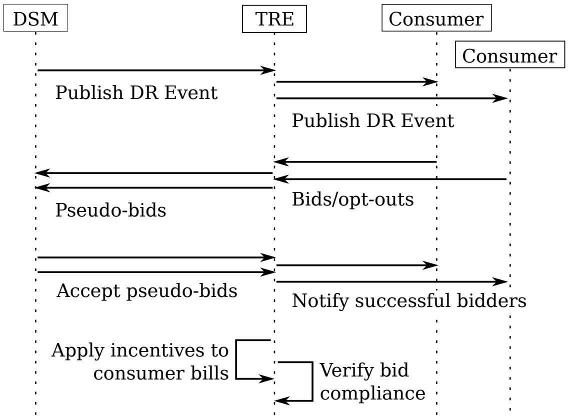

Fig. 1. Hybrid spatial and temporal aggregation for privacy-preserving DR.

F. Alternative Approaches

Previous work has used secure MPC to achieve the privacy objectives in the monitoring information flow [7]. Fundamentally, MPC aims to calculate an overall result from a set of private inputs. Whilst this could be used for either of the uni-directional information flows, it could not be used directly in the bi-directional DR information flow. Demand bidding would require a new type of MPC that allows the DSM to selectively accept bids and privately communicate different responses to individual consumers rather than computing an overall solution. Investigation of the feasibility of this alternative approach is beyond the scope of this paper.

G. Implementation Considerations

To ensure availability, it is expected that a geographically distributed network or TREs will be used. If a TRE fails, the affected consumers will re-connect to a different TRE since the aggregation groups are dynamically defined. Given the simplicity of the processing it performs, the TRE is not expected to add significant communication latency and since our architecture does not increase the number of messages sent by consumers it will not increase the network load. Since no modifications are required on the consumer side, this architecture can be deployed incrementally in parallel with the roll-out of smart meters.

IV. Establishing Trust

In the most extreme case, consumers and service providers could be mutually distrusting. Consumers do not trust service providers with their fine-grained energy measurements [1] and service providers do not trust consumers to aggregate their own measurements or calculate their own bills. The key feature of our communication architecture is that the TRE is singularly trusted by these mutually distrusting entities. Unlike a trusted third party which is blindly trusted, the TRE uses Trusted Computing (TC) technologies and approaches to allow relying parties to verify its trustworthiness.

We focus on TC technologies standardized by the Trusted Computing Group (TCG). These make use of the Trusted Platform Module (TPM), a cryptographic co-processor securely integrated into the platform that provides isolated storage for cryptographic keys and a special set of Platform Configuration Registers (PCRs). Each PCR stores an integrity measurement in the form of cryptographic hash that cannot be directly written but can be *extended* with a new hash which is concatenated with the existing value and the hash of the result stored in the PCR. The PCRs provide integrity-protection for the log of all software that has been executed on the platform. Before any software is executed, the preceding software takes a hash of the new binary, adds it to the log and extends it into the PCRs. *Remote attestation* can then be used to prove the state of the platform to a remote verifier by sending a TPM-signed *quote* of the PCR values to the verifier.

Although TPMs are readily available, the use of remote attestation has been very limited due to the size and complexity of the software on modern general-purpose systems. Studies have shown that remote attestation of a typical web service involves approximately 300 integrity measurements with about 35 new measurements each month due to software updates [27]. Furthermore, a general-purpose system usually includes an operating system (OS) kernel consisting of millions of lines of code. Even if the software can be unambiguously identified, this complexity makes it infeasible for a remote party to make informed trust decisions.

In our architecture, the TRE avoids these scalability issues since it is a highly specialized system with a single well-defined unchanging purpose. The TRE has a minimal Trusted Computing Base (TCB) consisting of a purpose-built network stack, a limited number of cryptographic primitives and the simple information processing procedures described above. As a single-purpose system, the TRE requires neither an OS nor the ability to execute any other software. This makes it feasible to use TCG *secure boot* so software binaries are only executed if their hashes are on a pre-defined white-list. During normal operation, the PCRs will therefore always reach the same value thus allowing sensitive information, such as cryptographic keys, to be *sealed* to this state. Most importantly, the TRE can use remote attestation to prove its state to all relying parties and due to its minimal TCB, this attestation can be used to make informed trust decisions. Since the TRE must be unambiguously identified, there is no need to use privacy-preserving attestation protocols. Each TRE uses a consistent Attestation Identity Key (AIK), endorsed by a regulatory authority, and a simple challenge-response attestation protocol:

$$\text{Verifier} \rightarrow \text{TRE:} \quad attestation\ request,\ n_t^v$$

$$\text{Verifier} \leftarrow \text{TRE:} \quad ML_t^{TRE},\ TS_t,\ \{n_t^v,\ PCR_t^{TRE}\}_{Sig(AIK)}$$

At time t, the verifier supplies nonce n_t^v. The TPM generates a signature over n_t^v and the current PCR values PCR_t^{TRE} using its AIK and sends this to the verifier with the current measurement list ML_t and a timestamp TS_t. The most significant performance constraint is the TPM's quote operation since current generation TPMs (version 1.2) are not usually designed to provide high throughput. To quantify this, we have performed micro-benchmarks on an Infineon TPM 1.2. The time taken by the TPM to perform one quote operation is approximately 731 milliseconds with a standard deviation of 0.7 ms over 3000 samples. Even with over 1000 consumers per TRE, this still allows every consumer to run the attestation protocol at least once every 15 minutes. It is anticipated that the TPM 2.0 will improve this performance.

V. Formal Analysis

The trust establishment procedure described above provides a technical basis for checking the exact system state of the TRE. However, in order to make a well-founded trusted decision, they must also be able to decide if this state provides the required security and privacy properties. Practically, these decisions are usually based on consistent experience of good behaviour but in some cases, formal methods can be used to analyse certain security and privacy properties resulting in much higher levels of assurance. Arguably the most critical aspects of the system are the communication protocols since a protocol flaw could have catastrophic effects. To mitigate against this risk, we have conducted an automated formal analysis of all the communication protocols in our architecture. This analysis is summarized in this section with the full details presented in the accompanying technical report[1].

The analysis was conducted using an enhanced version of the Casper/FDR security protocol analysis tool [9]. Given an abstract description of a communication protocol, Casper/FDR automatically checks claims about the security properties of secrecy and authentication. For this research, we augmented the tool to model and automatically analyse the privacy properties of undetectability and unlinkability. This tool has identified potential flaws in the security and/or privacy properties of several smart grid communication protocols from recent literature including demand bidding protocols. Although the tool cannot prove that a particular protocol is secure or privacy-preserving in an absolute sense, it can determine whether certain properties hold for a specific adversary model. For this analysis we have therefore used the adversary model for bi-directional smart grid communication defined in [2]. Table I shows the desired security and privacy properties of our communication architecture. We have used the enhanced Casper/FDR tool to perform a systematic analysis of these properties for each of the protocols in our architecture. This analysis has shown that none of the flaws identified in other protocols are present and we can therefore claim that, with respect to the defined adversary model, our architecture achieves its security and privacy objectives.

[1] https://www.cs.ox.ac.uk/people/andrew.paverd/tre/tre-smart-grid.pdf

TABLE I
SECURITY AND PRIVACY PROPERTIES ANALYSED IN OUR ARCHITECTURE

Security properties:

- Only authorized consumers can submit consumption measurements and DR bids *[authentication]*.
- Consumers cannot submit multiple measurements in a single period *[authentication]*.
- Unauthorized modifications of measurements or bids are detected *[integrity]*.
- Consumers cannot impersonate each other *[authentication]*.

Privacy properties:

- Measurements and bids cannot be viewed by external adversaries *[confidentiality]*.
- Only the TRE can detect if a specific consumer has placed a DR bid *[undetectability]*.
- Measurements, bids and DR incentives can only be linked to individual consumers by the TRE *[unlinkability]*.

VI. CONCLUSION

Due to the requirement of full bi-directional communication between service providers and individual consumers, current techniques for enhancing privacy in smart metering cannot be used address the privacy concerns arising from incentive-based DR protocols such as demand bidding. To address this challenge, we have proposed a unified communication architecture combining all three major information flows and making use of a Trustworthy Remote Entity (TRE). We have shown how the TRE facilitates demand bidding whilst preserving consumers' privacy using a combination of spatial and temporal aggregation and differential privacy. The TRE is fundamentally different from a trusted third party in that it uses Trusted Computing technologies (secure boot and remote attestation) to provide a technical mechanism for establishing its trustworthiness. Preliminary micro-benchmarks on a current generation TPM confirm the feasibility of this approach from a performance perspective. We have also conducted a systematic formal analysis of the communication protocols in our architecture using an automated analysis tool. This has shown that our architecture achieves its security and privacy objectives with respect to the defined adversary model.

Acknowledgements: This research was conducted as part of the *Future Home Networks and Services* project at the University of Oxford, funded by BT.

REFERENCES

[1] C. Cuijpers and B.-J. Koops, "Smart Metering and Privacy in Europe: Lessons from the Dutch Case," Feb. 2013.

[2] A. Paverd, A. Martin, and I. Brown, "Security and Privacy in Smart Grid Demand Response Systems," in *Second Open EIT ICT Labs Workshop on Smart Grid Security - SmartGridSec14*, 2014.

[3] United States Department of Energy, "Benefits of Demand Reponse in Electricity Markets and Recommendations for Achieving Them," Tech. Rep. February, 2006. [Online]. Available: http://energy.gov/oe/downloads/benefits-demand-response-electricity-markets-and-recommendations-achieving-them-report

[4] M. Albadi and E. El-Saadany, "A summary of demand response in electricity markets," *Electric Power Systems Research*, vol. 78, no. 11, Nov. 2008.

[5] OASIS, "Energy Interoperation Version 1.0," T. Considine, Ed., 2013.

[6] M. Karwe and J. Strüker, "Maintaining Privacy in Data Rich Demand Response Applications," in *First Open EIT ICT Labs Workshop on Smart Grid Security - SmartGridSec12*, 2013.

[7] C. Rottondi and G. Verticale, "Privacy-friendly appliance load scheduling in smart grids," in *Fourth IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Oct. 2013.

[8] M. Jawurek, M. Johns, and K. Rieck, "Smart metering de-pseudonymization," in *Proc. Computer Security Applications Conference - ACSAC '11*, 2011.

[9] A. J. Paverd, A. Martin, and I. Brown, "Modelling and Automatically Analysing Privacy Properties for Honest-but-Curious Adversaries," Tech. Rep., 2014. [Online]. Available: https://www.cs.ox.ac.uk/people/andrew.paverd/casper/casper-privacy-report.pdf

[10] C. Efthymiou and G. Kalogridis, "Smart Grid Privacy via Anonymization of Smart Metering Data," in *First IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2010.

[11] F. Borges, L. A. Martucci, and M. Muhlhauser, "Analysis of privacy-enhancing protocols based on anonymity networks," in *Third IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Nov. 2012.

[12] C. Rottondi, G. Mauri, and G. Verticale, "A data pseudonymization protocol for Smart Grids," in *2012 IEEE Online Conference on Green Communications (GreenCom)*, Sep. 2012.

[13] M. Stegelmann and D. Kesdogan, "GridPriv: A Smart Metering Architecture Offering k-Anonymity," in *Proc. IEEE International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, Jun. 2012.

[14] F. D. Garcia and B. Jacobs, "Privacy-Friendly Energy-Metering via Homomorphic Encryption," in *Proceedings of the 6th international conference on Security and trust management*, 2011.

[15] F. Li, B. Luo, and P. Liu, "Secure Information Aggregation for Smart Grids Using Homomorphic Encryption," in *First IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2010.

[16] J.-M. Bohli, C. Sorge, and O. Ugus, "A Privacy Model for Smart Metering," in *2010 IEEE International Conference on Communications Workshops*, May 2010.

[17] G. Ács and C. Castelluccia, "I have a DREAM!: differentially private smart metering," in *Proceedings of the 13th international conference on Information hiding - IH'11*, 2011.

[18] K. Kursawe, G. Danezis, and M. Kohlweiss, "Privacy-friendly aggregation for the smart-grid," in *Proceedings of the 11th international conference on Privacy enhancing technologies - PETS'11*, 2011.

[19] M. Jawurek, M. Johns, and F. Kerschbaum, "Plug-in privacy for smart metering billing," in *Proceedings of the 11th international conference on Privacy enhancing technologies*, 2011.

[20] A. Rial and G. Danezis, "Privacy-preserving smart metering," in *Proceedings of the 10th annual ACM workshop on Privacy in the electronic society - WPES '11*, Oct. 2011.

[21] G. Danezis, M. Kohlweiss, and A. Rial, "Differentially private billing with rebates," in *Proceedings of the 13th international conference on Information hiding - IH'11*, 2011.

[22] C. Dwork, "Differential Privacy," in *Automata, Languages and Programming*, ser. Lecture Notes in Computer Science, M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, Eds., 2006, vol. 4052.

[23] A. Bartoli, J. Herna andndez Serrano, M. Soriano, M. Dohler, A. Kountouris, and D. Barthel, "Secure Lossless Aggregation for Smart Grid M2M Networks," in *First IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2010.

[24] R. Petrlic, "A privacy-preserving Concept for Smart Grids," in *Sicherheit in vernetzten Systemen: 18. DFN Workshop*, 2010.

[25] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating Noise to Sensitivity in Private Data Analysis," in *Theory of Cryptography*, ser. Lecture Notes in Computer Science, S. Halevi and T. Rabin, Eds. Springer Berlin Heidelberg, 2006, vol. 3876.

[26] T. de Souza, J. Wright, P. O'Hanlon, and I. Brown, "Set Difference Attacks in Wireless Sensor Networks," in *Security and Privacy in Communication Networks*, ser. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, A. Keromytis and R. Pietro, Eds., 2013, vol. 106.

[27] J. Lyle and A. Martin, "Engineering attestable services," in *Proceedings of the 3rd international conference on Trust and trustworthy computing - TRUST '10*, Jun. 2010.