Consumer privacy protection using flexible thermal loads: Theoretical limits and practical considerations

Applied Energy 281 (2021) 116075

Contents lists available at ScienceDirect

Applied Energy

journal homepage: www.elsevier.com/locate/apenergy

Consumer privacy protection using flexible thermal loads: Theoretical limits and practical considerations

Jun-Xing Chin [a,*], Kyri Baker [b], Gabriela Hug [c]

[a] Singapore-ETH Centre, ETH Zurich, Singapore
[b] Architectural Engineering, University of Colorado Boulder, CO, USA
[c] Power Systems Laboratory, ETH Zurich, Zurich, Switzerland

ARTICLE INFO

Keywords:
Consumer privacy
Energy management
Energy storage
Flexible thermal loads
Smart meter

ABSTRACT

The increasing adoption of smart meters introduces growing concerns about consumer privacy risks stemming from high resolution metering data. To counter these risks, there have been various works in actively shaping the grid-visible energy consumption profile using controllable loads such as energy storage systems (ESSs) and flexible consumer loads. In this paper, we compare the use of flexible thermal-based consumer loads (FTLs) against ESSs for consumer privacy protection. By first assuming ideal conditions, and subsequently bringing them closer to reality, the limitations of using FTLs for privacy protection are identified. Through theoretical analyses and realistic simulations, it is shown that, due to the limitations in the operation of FTLs, without significant over-sizing of systems and sacrifices in consumer comfort, FTLs of much higher equivalent energy storage capacity are required to afford the same level of protection as ESSs. Nonetheless, given their increasing ubiquity, controllable FTLs should be considered for use in consumer privacy protection.

1. Introduction

Spurred by grid modernisation efforts, the adoption rate of advanced metering infrastructure (AMI) using smart meters (SMs) has risen steadily across the globe in recent years. On one hand, this enables the development of efficient data-driven grid operation and management methods [1]. On the other hand, the high-frequency measurement data provided by the AMI can be used to derive private information of consumers, such as their lifestyle habits, occupation, and religious inclinations [1–3]. The authors in [1] provide a comprehensive overview of applications (and information) that can be derived from SM data, while in [3], the authors explore the granularity of SM measurements required to infer specific household activities, and show that some private information can still be inferred at an hourly resolution. More importantly, the authors of [2] find that the existing laws in the US are unclear regarding customer energy data usage, which potentially paves the way for its exploitation. Moreover, a 2017 survey in the US has shown that utilities pose high privacy risks, and are not highly trusted by consumers [4]. Even in the presence of clear laws that prevent the exploitation of SM data by utility companies, such as the European Union's General Data Protection Regulation [5], the underlying metering infrastructure is still vulnerable to cyber-attacks, which may lead to SM data disclosure to malicious adversaries [6].

This has led to concerns regarding privacy risks [7], and push-backs against the use of SMs, delaying and potentially altering the scope of their deployment, e.g., in the Netherlands [8]. These concerns have motivated works in quantifying and mitigating these risks, such as [9–17]. Nonetheless, quantifying privacy in a meaningful manner remains an ongoing research challenge [17]. In [15] and [16], the authors use the performance of specific data analytics applications, *i.e.*, non-intrusive load monitoring (NILM) and socio-demographic classifiers, respectively, as measures of consumer privacy. While they allow simple and meaningful interpretation of a consumer's level of privacy, NILM techniques, which are reliant on appliance load signatures, are very sensitive to small perturbations in the data [15]; whereas machine learning-based classifiers are sensitive to rudimentary privacy protection measures [16]. This makes them less robust as a measure of privacy against more sophisticated adversaries. On the other hand, information theoretic privacy measures, such as mutual information (MI) [9] and differential privacy [14] are attack-agnostic and offer more robust privacy guarantees, but are less readily interpretable in a meaningful manner [17].

A recent article by Giaconi et al. [18] provides a high-level overview of privacy protection methods for consumers with SMs. In particular, privacy protection schemes can be categorised into two main families, namely *smart meter data manipulation* (SMDM) schemes, and *user*

* Corresponding author.
E-mail addresses: chin@eeh.ee.ethz.ch (J.-X. Chin), Kyri.Baker@colorado.edu (K. Baker), hug@eeh.ee.ethz.ch (G. Hug).

https://doi.org/10.1016/j.apenergy.2020.116075
Received 13 March 2020; Received in revised form 10 September 2020; Accepted 19 October 2020
Available online 27 October 2020

demand shaping (UDS) schemes [18]. SMDM schemes modify the SM data before it is transmitted, and include aggregating SM measurements before transmission to the data collector [19–21]; anonymising the SM measurements to decouple SM data from individual households [22,23]; and the differential privacy-based addition of noise [24–26]. However, these methods require trusted third parties, either in the processing of the data, or in the supply and installation of SMs with privacy-preserving firmware. UDS methods, on the other hand, physically alter the physical energy consumption profiles of consumers recorded by the SMs (*grid load*), such that they no longer reveal the private information contained in the underlying privacy-sensitive consumer load profiles (*sensitive load*). This is achieved by actively controlling loads to shape the grid load profile, ideally decoupling it from the sensitive load profile. UDS methods can typically be implemented behind-the-meter, which avoids the need for a trusted third-party.

UDS methods can be further classified into those using energy storage systems (ESSs), those controlling flexible consumer loads, and those using a combination of the two. Fig. 1 illustrates a possible system setup for UDS methods, which is governed by the equation:

$$Y = X + S. \tag{1}$$

Hence, the flexibility of controllable loads S, such as ESSs and flexible consumer loads, is used to influence what can be derived from the grid load Y about the sensitive load X.

There are numerous recent UDS schemes that only use ESSs (also known as battery load hiding), e.g., load levelling [10], limiting the load profile to distinct steps [11], and directly minimising an approximate of MI [13]. In [27], the authors derive theoretical privacy guarantees for consumers with ESSs and renewable energy sources based on ideal assumptions, and show that, while it is possible to numerically evaluate the privacy bounds for realistic batteries using the Blahut–Arimoto algorithm, it is computationally intractable in practice. Arzamasov et al. provide a more recent overview of SM related privacy measures for ESS-based UDS methods in their recent work [17], where they also found that the choice of privacy metrics and the characteristics of a consumer's load profile greatly affect the relative performance of ESS-based UDS schemes. They argue that an ideal privacy measure would be the reconstructability of the original unprotected consumer load profile. However, assessing the reconstructability of the consumer load profile given a specific privacy protection scheme is a non-trivial problem that remains to be solved.

On the other hand, UDS methods utilising flexible consumer loads are scarce in the literature. One such UDS scheme, proposed in [28], utilises the flexible consumer loads to hide occupancy by using artificial signature injection and partial load flattening. The authors then verify their scheme by testing the resultant load profiles using a few occupancy detection algorithms. Another flexible consumer load-based UDS scheme is given in [29], where the authors use flexible consumer loads aided with batteries for privacy protection. However, no in-depth assessment or discussion on the performance of the proposed scheme is included. In [15], optimised electric vehicle charging and an electric furnace are used to obscure recoverable information from NILM techniques. Notwithstanding, the use of flexible consumer loads for general privacy protection irrespective of the adversarial model, and their performance against schemes based on ESSs, are not well studied.

With the development of grid communications infrastructure and the proliferation of smart appliances, there are also considerable advances in home energy management systems (HEMSs) that enable the coordination and scheduling of home appliances. HEMSs allow for the optimisation of residential electricity consumption patterns in order to improve efficiency, economics, and the reliability of residential buildings with regards to their role in the grid and occupant comfort [30]. Given increasing interest in HEMSs and the ubiquity of flexible consumer loads, this paper explores the use of HEMS-controlled flexible consumer loads in order to mask the private information contained in the grid load about the sensitive load. Specifically, the contributions of this paper are three-fold:

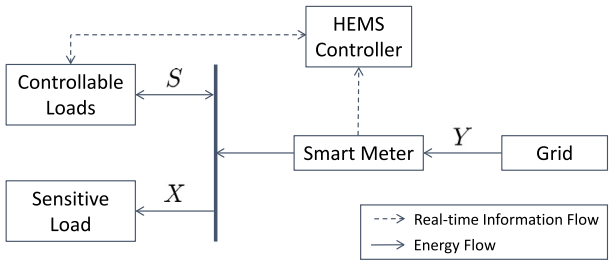

Fig. 1. Possible UDS system setup, where the controllable load is controlled by the HEMS controller to actively shape the grid visible load to mask the private information in the privacy-sensitive load.

- The concept of privacy and flexible consumer loads is formalised for households with smart meters.
- The theoretical limits of privacy protection using ESSs and flexible thermal-based consumer loads are analysed, with the findings validated for real-life applications using realistic numerical simulations.
- To the best of our knowledge, this paper is the first work to directly compare consumer privacy protection of systems using ESSs against those using flexible thermal-based consumer loads.

The rest of this paper is structured as follows: Section 2 provides a brief overview of quantifying privacy loss for consumers with smart meters; Section 3 briefly discusses the use of flexible consumer loads for privacy protection; Section 4 provides an analytical comparison between consumer privacy protection using ESSs and flexible thermal-based consumer loads; Section 5 details the controller design of a HEMS for comparison of realistic systems; Section 6 presents numerical results; and Section 7 concludes the paper.

2. Quantifying consumer privacy loss

As previously mentioned, one measure of consumer privacy loss is the mutual information between the sensitive load X and the grid load Y [9,12], which measures the amount of information Y reveals about X and vice versa. Mutual information is capable of modelling nonlinear relationships between variables, unlike using correlation coefficients, for example. The MI between X and Y, which are random processes, can be given as the average MI between the random variables X_τ and Y_τ that make up the processes [12,31], *i.e.*,

$$I(X;Y) = \frac{1}{k}\sum_{\tau=1}^{k} I(X_\tau;Y_\tau)\ , \tag{2}$$

where $I(X_\tau;Y_\tau)$ is the MI between the random variables X_τ and Y_τ, and k is the number of random variable pairs. This concept of average MI will be used in Section 4 for the analysis of consumer privacy protection.

Given two random variables X_τ and Y_τ, the MI between them is given by a function of their joint probability distribution function (PDF) p_{X_τ,Y_τ}, and marginal distributions, p_{X_τ}, and p_{Y_τ}. These PDFs are typically unknown, and must be estimated. Assuming that multiple samples of X_τ and Y_τ are available, the PDFs can be estimated using the histogram method. Hence, only for the purpose of estimating these PDFs, assume that the protected and grid loads have finite support, *i.e.*, $X_\tau \in \mathcal{X}_\tau := \{\bar{x}^1, \bar{x}^2, \dots, \bar{x}^m\}$, and $Y_\tau \in \mathcal{Y}_\tau := \{\bar{y}^1, \bar{y}^2, \dots, \bar{y}^n\}$. Then, the MI between X_τ and Y_τ can be given as

$$I(X_\tau;Y_\tau) := \sum_{i=1}^{m}\sum_{j=1}^{n} p_{X_\tau,Y_\tau}(\bar{x}^i,\bar{y}^j)\log\frac{p_{X_\tau,Y_\tau}(\bar{x}^i,\bar{y}^j)}{p_{X_\tau}(\bar{x}^i)p_{Y_\tau}(\bar{y}^j)}, \tag{3}$$

where $p_A(a)$ denotes the probability of $A = a$, and log is the base-2 logarithm. For the rest of the paper, we further denote the realisations of the random variables with lowercase letters, A^{mean} as the average

value of A, $A^{\min}$ as the minimum value that A can take, and $A^{\max}$ as the maximum value of A.

As X and Y are continuous in reality, the PDF estimates become more accurate with an increase in m and n; but this also requires more samples to prevent over-fitting. It follows that in order to minimise leakage of privacy-sensitive information, one needs to minimise the MI between the sensitive and grid loads. This can be done either through UDS or SMDM methods as described in Section 1; and for UDS methods, using either ESSs, flexible consumer loads, or a combination thereof.

3. Flexible consumer loads and consumer privacy

The term "flexible consumer loads" include thermal loads such as hot water heaters and space conditioning, schedulable loads such as clothes and dishwashers, and interruptible loads such as the charging of electric vehicles. From a privacy perspective, flexible consumer loads can broadly be classified into the following categories:

(a) Flexible consumer loads that are not privacy-sensitive, *i.e.*, their usage does not reveal privacy-sensitive information about the consumer, nor are their presence in a household considered sensitive private information; e.g., electric space heaters within a house with high thermal inertia, in a community where their presence is the norm.
(b) Flexible consumer loads that are privacy-sensitive with regards to their time-of-use, but not their presence in the household; e.g., electric stoves in a community where their presence is the norm.
(c) Flexible consumer loads that are privacy-sensitive, *i.e.*, both their time-of-use and presence in a household reveal sensitive private information; e.g., electric stoves in a community where households typically cook with gas stoves.

When controlling flexible consumer loads to shape a user's demand and reduce their information leakage, the privacy sensitivity of the loads themselves need to be considered. There are no privacy issues arising from their usage if the flexible consumer loads are of the first category. For loads of the second category, using them to mask the sensitive load inherently also masks the private information they reveal: their time-of-use is shifted and thus, the private information revealed by their original time-of-use is masked. However, if the flexible consumer loads are of the third category, then the privacy-protection problem also needs to consider whether the resulting grid load is able to mask the electrical signature of the flexible consumer loads, *i.e.*, whether the sensitive load is able to sufficiently distort the signatures of the flexible consumer loads as well.

To simplify the analysis, we consider the use of flexible consumer loads within the first two categories in UDS privacy-protection schemes in this paper. Moreover, we limit our analysis to flexible thermal-based consumer loads (flexible thermal loads or FTLs) due to their ability to 'store' thermal energy, and are more likely to be interruptible compared to other types of flexible consumer loads, such as washing machines that have minimum cycle times. Inductive FTLs, such as heat pumps, have complex on/off cycles and electrical signatures, making the analysis of their effectiveness in privacy protection complicated. Hence, in order to draw meaningful conclusions, we will focus on resistance-based FTLs, such as electric-resistance water heaters, and electric-resistance space heaters. In the next section, we will compare the theoretical performance of ESS-based UDS schemes against those using resistance-based FTLs.

4. Comparing privacy protection using energy storage systems and flexible thermal loads

Setting aside the distinctive constraints of both ESSs and FTLs, the privacy protection afforded by them for UDS differs in one key aspect: ESSs are able to both charge and discharge, *i.e.*, increase or decrease grid load; while traditional residential FTLs are only able to 'charge', *i.e.*, they can draw power from the grid, but typically cannot provide power back to the grid.

Let $H(\cdot) := -\sum p(\cdot)\log p(\cdot)$ be the Shannon entropy function, with $p(\cdot)$ being the probability of the variable and $H(\cdot)$ being minimal when the outcome is certain, and maximal when the underlying distribution is uniform. Additionally, assume that the following is true:

(a) No energy wastage is permitted.
(b) The power ratings of the ESS and FTL are sufficiently large to compensate for the difference between the maximum and minimum consumer load, *i.e.*, $P_{\text{ess}}^{\max}, P_{\text{th}}^{\max} \geq X^{\max} - X^{\min}$.
(c) The controller has perfect knowledge of the efficiency curves of the ESS and FTL.
(d) The controller has perfect knowledge of the consumer load X and its average X^{mean}.
(e) The ESS has infinite energy storage capacity.
(f) Either the FTL has infinite thermal storage capacity, or it holds, for the average electrical equivalent of the consumer thermal demand $D_{\text{th}}^{\text{mean}}$, that $D_{\text{th}}^{\text{mean}} \geq P_{\text{th}}^{\max}$.
(g) The FTL demand is continuous, *i.e.*, it is not a step-load.
(h) Both ESS and FTL have an initial state-of-charge of 0.5.

Using MI as the measure of privacy, the differences in achievable privacy protection by both technologies are discussed in the remainder of this section.

4.1. The loads are independent and identically distributed

Let the random variable pair (X, Y), and its marginals X and Y be independent and identically distributed (i.i.d.). Then, by definition, MI, $I_{\text{iid}}(X;Y)$ can also be written as a function of their Shannon entropies,

$$I_{\text{iid}}(X;Y) = H(X) + H(Y) - H(X,Y) \tag{4}$$

$$= H(Y) - H(Y|X)\ . \tag{5}$$

From (4), it is trivial to see that $I_{\text{iid}}(X;Y)$ can be minimised by either maximising $H(X,Y)$, assuming $H(X,Y)$ increases at the same or higher rate than $H(Y)$; or by minimising $H(Y)$, assuming $H(Y)$ decreases at the same or higher rate than $H(X,Y)$. Note that as $H(X)$ is fixed, $H(Y)$ and $H(X,Y)$ are affected similarly (either both increase or both decrease) by a given control action. Moreover, we have the following propositions:

Proposition 1. *$I_{iid}(X;Y)$ is minimal when $H(Y)$ is minimal, i.e., when $|\mathcal{Y}' := \{y \in \mathcal{Y} \mid p_Y(y) > 0\}|$ is minimal.*

Proof. The distribution of the sensitive load $p_X(x)$ is uncontrollable, non-uniform, and the number of outcomes with non-zero probability is non-singular, *i.e.*, $|\mathcal{X}' := \{x \in \mathcal{X} \mid p_X(x) > 0\}| > 1$. Hence, $H(X)$ is greater than zero. Since $p_{X,Y}(x,y)$ is non-uniform, as $p_X(x)$ is non-uniform, $H(X,Y)$ is limited by the given $p_X(x)$. Therefore, it follows that $I_{\text{iid}}(X;Y)$ is minimal when $H(Y)$ is minimal, *i.e.*, when $|\mathcal{Y}'|$ is minimal, where $\mathcal{Y}' := \{y \in \mathcal{Y} \mid p_Y(y) > 0\}$, instead of when $H(X,Y)$ is maximal. □

Proposition 2. *$I_{iid}(X;Y)$ is minimal when $H(Y|X)$ is maximal.*

Proof. Given a fixed and non-uniform sensitive load distribution $p_X(x)$, $H(Y|X)$ is maximal when $p_{Y|X}(y|x)$ are uniform distributions for each value of x. Since $p_Y(y) = \sum_{\mathcal{X}} p_{X,Y}(x,y) = \sum_{\mathcal{X}} p_{Y|X}(y|x)p_X(x)$, it is also a uniform distribution when $p_{Y|X}(y|x)$ are uniform distributions for each value of x. Therefore, $H(Y) = H(Y|X)$ when $H(Y|X)$ is maximal, and $I_{\text{iid}}(X;Y) = 0$, which is its minimal value. □

As FTLs cannot 'discharge' (reduce the grid load), and therefore, cannot achieve a uniform distribution for $p_{Y|X}(y|x)$, the following analysis is based on Proposition 1. It is trivial to see that perfect privacy, $I_{\text{iid}}(X;Y)=0$ can be achieved by maintaining a constant grid load, y^*, where $p_Y(y^*)=1$,and $p_Y(y)=0\ \forall y\neq y^*$. Let the grid load achieved using the ESS be denoted by Y_{ess} and that of FTL by Y_{th}, then there exists y^*_{ess} and y^*_{th} such that $I_{\text{iid}}(X;Y^*_{\text{ess}})=I_{\text{iid}}(X;Y^*_{\text{th}})=0$. While y^*_{ess} can be any arbitrary value $Y^{\min}\leq y^*_{\text{ess}}\leq Y^{\max}$, there is less flexibility for y^*_{th}, with $y^*_{\text{th}}\geq X^{\max}$. Nonetheless, the theoretical maximum privacy can be achieved by both technologies given ideal assumptions.

In reality, storage capacity is finite, and for most consumers, it would be unreasonable to assume that the system is undersized, *i.e*, $D^{\text{mean}}_{\text{th}}\geq P^{\max}_{\text{th}}$, where $D^{\text{mean}}_{\text{th}}$ is the electrical equivalent of the average power consumption required in order to maintain consumer comfort. Therefore, assumptions (e) and (f) are made more realistic such that the storage capacity is finite, but sufficiently large to average out consumer load (or thermal demand) over a finite period of time. Additionally, average thermal demand is now assumed to be large, but less than the FTL power rating and that $D^{\text{mean}}_{\text{th}}+X^{\text{mean}}<X^{\max}$. For ESSs, the controller would now need to select a constant grid load such that $y^*_{\text{ess}}=X^{\text{mean}}+l_{\text{ess}}$, where l_{ess} is the round trip loss of the ESS. This allows a constant y^*_{ess} that does not empty or fully charge the ESS. As it would be possible to sustain y^*_{ess} indefinitely, $I_{\text{iid}}(X;Y_{\text{ess}})=I_{\text{iid}}(X;Y^*_{\text{ess}})=0$. For FTLs, it follows that $y^*_{\text{th}}=D^{\text{mean}}_{\text{th}}+X^{\text{mean}}$, and that $y^*_{\text{th}}<X^{\max}$ in most realistic cases. Assume that $I_{\text{iid}}(X;Y_{\text{th}})$ is still minimised by actuating y^*_{th} whenever possible. In this case, we now also have $y=x\neq y^*_{\text{th}}\ ,\forall x>y^*_{\text{th}}$. Let k be the total number of samples, and $g(k)$ be the number of instances where $x>y^*_{\text{th}}$, then

$$I_{\text{iid}}(X;Y_{\text{th}})=\frac{k-g(k)}{k}I(X;Y^*_{\text{th}})+\frac{g(k)}{k}H(X) \tag{6a}$$

$$=\frac{g(k)}{k}H(X)\ , \tag{6b}$$

where (6b) follows from the fact that $I_{\text{iid}}(X;Y^*_{\text{th}})=0$ and $y=x\ ,\forall x>y^*_{\text{th}}$. Thus, $I_{\text{iid}}(X;Y_{\text{ess}})<I_{\text{iid}}(X;Y_{\text{th}})$ as $H(X)>0$, *i.e.*, privacy loss using FTLs for UDS schemes is, under the given assumptions on equivalent storage size, greater than those using ESSs given these assumptions.

In theory, MI can also be minimised by maximising $H(Y|X)$ if $|\mathcal{Y}':=\{y\in\mathcal{Y}\mid p_Y(y)>0\}|$ is too large, *i.e*, when there are too many different values of $x>y^*_{\text{th}}$. However, given that FTLs cannot 'discharge', $p_{Y|X}(y|x)$ cannot be uniform distributions. Therefore, $H(Y)-H(Y|X)>0$, as $H(Y)\neq H(Y|X)$ in this case, again, resulting in $I_{\text{iid}}(X;Y^*_{\text{th}})>I_{\text{iid}}(X;Y^*_{\text{ess}})$.

4.2. The loads are first-order Markov processes

The random variables (X,Y), X, and Y are not i.i.d. in reality, and could be better modelled using first-order Markov processes [32], of which the MI, $I_{\text{m}}(X;Y)$ [12] is given by

$$I_{\text{m}}(X;Y)=\frac{1}{k}\left[\sum_{\tau=2}^{k}I(X_\tau,X_{\tau-1};Y_\tau,Y_{\tau-1})-\sum_{\tau=3}^{k}I(X_{\tau-1},Y_{\tau-1})\right]\ . \tag{7}$$

Expressing (7) in terms of entropy,

$$\begin{aligned}I_{\text{m}}(X;Y)=\frac{1}{k}\Bigg\{&H(X_2,X_1)+H(Y_2,Y_1)-\\&H(X_2,X_1,Y_2,Y_1)+\sum_{\tau=3}^{k}\Big[H(X_\tau,X_{\tau-1})+\\&H(Y_\tau,Y_{\tau-1})-H(X_\tau,X_{\tau-1},Y_\tau,Y_{\tau-1})-\\&H(X_{\tau-1})-H(Y_{\tau-1})+H(X_{\tau-1},Y_{\tau-1})\Big]\Bigg\}\ .\end{aligned}$$

Note that if the random variables (X,Y), X, and Y are higher-order Markov processes, then (7) forms the upper bound on the actual MI [12]. As $k\to\infty$,

$$\begin{aligned}I_{\text{m}}(X;Y)\approx\frac{1}{k}\Bigg\{\sum_{\tau=3}^{k}\Big[&H(X_\tau,X_{\tau-1})+H(Y_\tau,Y_{\tau-1})-\\&H(X_\tau,X_{\tau-1},Y_\tau,Y_{\tau-1})-H(X_{\tau-1})-\\&H(Y_{\tau-1})+H(X_{\tau-1},Y_{\tau-1})\Big]\Bigg\}\ .\end{aligned}$$

It is trivial to see that Proposition 1 still holds, and that $I_{\text{m}}(X;Y)$ is minimal when the entropy of Y is minimal. Moreover, when assumptions (a) to (g) hold, then both $I_{\text{m}}(X;Y^*_{\text{ess}})$ and $I_{\text{m}}(X;Y^*_{\text{th}})$ are minimal and equal to zero. Now, assume that the Markov processes (X,Y), X, and Y are also stationary, *i.e.*, $H(X_1)=H(X_2)=\cdots=H(X_k)$, $H(Y_1)=H(Y_2)=\cdots=H(Y_k)$, $H(X_1,X_2)=H(X_2,X_3)=\cdots=H(X_{k-1},X_k)$, $H(Y_1,Y_2)=H(Y_2,Y_3)=\cdots=H(Y_{k-1},Y_k)$, and that assumptions (e) and (f) are made more realistic as in the i.i.d. case. Then, $I_{\text{m}}(X;Y_{\text{ess}})=I_{\text{m}}(X;Y^*_{\text{ess}})=0$, while

$$\begin{aligned}I_{\text{m}}(X;Y_{\text{th}})\approx&\ \frac{g_1(k)}{k}\cdot 0+\frac{g_2(k)}{k}H(X_\tau)+\\&\frac{g_3(k)}{k}\Big[H(X_\tau,X_{\tau-1})-H(X_\tau)\Big]\\=&\ \frac{g_2(k)-g_3(k)}{k}H(X_\tau)+\\&\frac{g_3(k)}{k}H(X_\tau,X_{\tau-1}),\quad \tau\in\{2,3,\dots,k\}\ ,\end{aligned}$$

where the function $g_1(k)$ gives the number of instances where $(y_{\text{th},\tau}=y_{\text{th},\tau-1}=y^*_{\text{th}})$ or $(y_{\text{th},\tau}=y^*_{\text{th}},\ y_{\text{th},\tau-1}=x_{\text{th},\tau-1})$, $g_2(k)$ is the number of instances where $(y_{\text{th},\tau}=x_{\text{th},\tau},\ y_{\text{th},\tau-1}=y^*_{\text{th}})$, $g_3(k)$ is the number of instances where $(y_{\text{th},\tau}=x_{\text{th},\tau},\ y_{\text{th},\tau-1}=x_{\text{th},\tau-1})$, and $g_1(k)+g_2(k)+g_3(k)=k-2$ [12]. As $H(X_\tau)>0$, $H(X_\tau,X_{\tau-1})>0$, and $H(X_\tau,X_{\tau-1})>H(X_\tau)$ (because X_τ and $X_{\tau-1}$ are not perfectly correlated), therefore, $I_{\text{m}}(X;Y_{\text{th}})>I_{\text{m}}(X;Y_{\text{ess}})$.

4.3. Privacy protection for actual systems

For actual systems, the load distributions vary according to the consumer household's state, and their characterisation is the subject of much research. Despite this, consumer privacy is protected if one can achieve a flat grid load that has zero entropy, *i.e.*, zero MI between the sensitive and grid loads. While assumptions (d) and (g) do not hold in reality, it would be possible to implement systems with sufficient storage capacity to average out consumer load (or thermal demand). For ESS-based schemes, one would be able to select y_{ess} close to y^*_{ess}, given a sufficiently large sample size, as the accuracy of the consumer load sample mean $\hat{X}^{\text{mean}}\to X^{\text{mean}}$ as $k\to\infty$. In addition to X^{mean}, the achievable privacy protection of FTL-based UDS schemes is also dependent on $D^{\text{mean}}_{\text{th}}$ and the ratio of $X^{\max}$ to X^{mean}, which are usually fixed and directly affect the number of instances when $y_{\text{th}}=y^*_{\text{th}}$. Note that a larger $X^{\max}$ to X^{mean} ratio would require a larger $D^{\text{mean}}_{\text{th}}$ to achieve the same level of privacy protection and vice versa. It would be difficult to compare the performance of actual ESS and FTL-based UDS privacy protection schemes, especially since there is a lot of uncertainty in the system parameters for FTLs. Even so, given the analysis above, the additional dependencies of FTL-based schemes (stochastic thermal demand and dependencies on the ambient environment), and the fact that most FTLs are step-loads, properly designed ESS-based schemes should outperform their FTL-based counterparts.

5. Formulation of the optimisation problem for numerical experiments

Our goal is to compare the performance of privacy protection using ESSs and FTLs in realistic systems to validate our findings from the

previous theoretical analysis by simulating a multi-objective model-predictive control-based HEMS controller. For FTLs, we analyse the use of electric hot water heaters (EWHs) and electric resistance space heaters (ERHs), as they better match the analysis in Section 4 compared to other FTL types. In this section, the modelling of the ESS and FTLs, the formulation of the privacy objective, and the overall optimisation problems used in the HEMS controllers are presented.

5.1. Privacy objective

There are numerous privacy measures in use across the many different privacy protection schemes available, as briefly summarised by the authors of [17,18]. To better match the analysis in Section 4, we adopt a privacy objective function that directly minimises an approximation of (3). This MI approximate, as proposed in [13], assumes that X and Y are i.i.d., and is given by:

$$\begin{aligned} I(X;Y) \approx \;& \tilde{I}(X_w;Y_w) \\ :=& \sum_{i=1}^{m}\sum_{j=1}^{n}\left(a_w^{ij}+\frac{1}{N_\varepsilon}\sum_{\tau=w}^{w+W} z_\tau^{ij}\right)\times \\ & \left\{\log\frac{a_w^{ij}}{b_w^j c_w^i}+\frac{\nu}{a_w^{ij}N_\varepsilon}\sum_{\tau=w}^{w+W} z_\tau^{ij}-\frac{\nu}{b_\tau^j N_\varepsilon}\sum_{\tau=w}^{w+W}\sum_{h=1}^{m} z_\tau^{hj}\right\}, \end{aligned} \tag{8}$$

at time w, where $W+1$ is the prediction horizon, a_w^{ij}, b_w^j, and c_w^i are constants used in the estimation of the PDFs $p_{X,Y}$, p_X and p_Y, N_ε is the total number of observations used in the estimate, including an additive smoothing constant, $\nu := 1/\log_e 2$, and $z_\tau^{ij} \in \{0,1\}$ are binary variables used to estimate the PDFs; see [13] for details on its derivation. While this MI approximate was shown to directly minimise the MI between X and Y, its scalability is limited by the number of binary variables z_τ^{ij}, which increases with the prediction horizon length and quantisation levels of X and Y. Hence, we relax binary variables z_τ^{ij}, *i.e.*, let $z_\tau^{ij} \in [0,1]$, in order to make (8) a convex function, and overcome the scalability issues identified in [13]. This relaxation affects the performance of the controller in terms of minimising MI, but this is outside the scope of this paper. The following constraints are required in the optimisation of (8):

$$\sum_{j=1}^{n} z_\tau^{i^* j} = 1 \tag{9}$$

$$z_\tau^{ij} = 0 \qquad ,\forall\, i \neq i^* \tag{10}$$

$$\sum_{j=1}^{n} z_\tau^{i^* j}\bar{y}^{j-1} \leq \; y_\tau < \sum_{j=1}^{n} z_\tau^{i^* j}\bar{y}^{j}\;, \tag{11}$$

where i^* is the index corresponding to the given value of x_τ, $\bar{y}^0 = Y^{\min}$, $\bar{y}^n = Y^{\max}$, and constraint (11) links the grid load to its PDF estimate and thus, the MI approximate.

5.2. Modelling of an ESS

Two variables, P_c and P_d, are used to model the instantaneous charging and discharging powers of the ESS, respectively, in order to capture the different loss factors during charge and discharge. Additionally, a binary variable B_{ess} is introduced to prevent the simultaneous charging and discharging of the ESS, necessitated by the fact that this is optimal at some time instances due to the privacy objective. While it would be ideal to have a realistic and convex ESS model, its derivation remains an ongoing area of research. Let E_τ be the energy remaining in the ESS at time τ. Then, the following constraints are used to model the ESS in the optimisation problem:

$$0 \leq P_{c,\tau} \leq B_{ess,\tau} P_c^{\max} \tag{12}$$

$$0 \leq P_{d,\tau} \leq (1 - B_{ess,\tau}) P_d^{\max} \tag{13}$$

$$0 \leq E_\tau \leq E^{\max} \tag{14}$$

$$E_{\tau+1} = E_\tau + \Delta_t(\eta_c P_{c,\tau} - \eta_d P_{d,\tau}) \tag{15}$$

$$S_\tau = P_{c,\tau} - P_{d,\tau} \qquad , \tag{16}$$

where η_c and η_d are the charging and discharging efficiencies of the ESS, respectively, and Δ_t is the interval of τ.

5.3. Modelling of an electric hot water heater

The thermodynamics in a hot water tank can be modelled by splitting the tank into several sections (nodes). A two-node EWH model proposed in [33] is adopted in order to better capture the thermodynamics of a real device. As the original model was developed for an electric heat pump, we modify it by replacing the coefficient of performance (COP) with one. Also, we assume a temperature deadband of 1°C around the temperature set-point. This water heater model is given by the following constraints in the optimisation problem:

$$\begin{aligned} T_{ewh,\tau+1}^{low} = T_{ewh,\tau}^{low} + \frac{\Delta_t}{C_{ewh}^{low}}\Big[& UA_{ewh}^{low}\left(T_{air,\tau}^{in} - T_{ewh,\tau}^{low}\right) + \\ & \Delta m_{hw,\tau} C_p\left(T_{ms} - T_{ewh,\tau}^{low}\right) + P_{ewh}^{\max} U_{ewh,\tau}^{low}\Big] \end{aligned} \tag{17}$$

$$\begin{aligned} T_{ewh,\tau+1}^{up} = T_{ewh,\tau}^{low} + \frac{\Delta_t}{C_{ewh}^{up}}\Big[& UA_{ewh}^{up}\left(T_{air,\tau}^{in} - T_{ewh,\tau}^{up}\right) + \\ & \Delta m_{hw,\tau} C_p\left(T_{ewh,\tau}^{low} - T_{ewh,\tau}^{up}\right) + P_{ewh}^{\max} U_{ewh,\tau}^{up}\Big] \end{aligned} \tag{18}$$

$$T_{ewh}^{absmin} \leq \; T_{ewh,\tau}^{up} \leq \; T_{ewh}^{absmax} \tag{19}$$

$$T_{ewh,\tau}^{low} \leq \; T_{ewh,\tau}^{up} \tag{20}$$

$$U_{ewh,\tau}^{low} + \; U_{ewh,\tau}^{up} \leq \; 1 \tag{21}$$

$$S_\tau = \Delta_t(P_{ewh}^{\max} U_{ewh,\tau}^{low} + P_{ewh}^{\max} U_{ewh,\tau}^{up}) \tag{22}$$

where superscripts *low* and *up* represent the values for the lower and upper nodes of the tank, respectively. $T_{ewh,\tau}$ is the water temperature of the node, $T_{air,\tau}^{in}$ is the indoor air temperature, $\Delta m_{hw,\tau}$ is the hot water draw, and $U_{ewh,\tau} \in [0,1]$ is the duty cycle of the EWH tank node at time τ. Also, C_{ewh} is the thermal capacitance of the tank node, UA_{ewh} is the heat loss coefficient of the node, C_p is the heat capacity of water, T_{ms} is the mains water temperature, $P_{ewh}^{\max}$ is the rated power of the EWH, T_{ewh}^{absmin} is the minimum water temperature required for safety (to mitigate Legionella bacterium growth in pipework), and T_{ewh}^{absmax} is the maximum permissible water temperature of the EWH. Furthermore, to take into account consumer comfort, variables $z_\tau^{comf} \in \mathbb{R}_{\geq 0}$ with constraints:

$$\left(T_{ewh}^{set} - 1°\mathrm{C}\right) - T_{ewh,\tau}^{low} \leq \; z_\tau^{comf} \tag{23}$$

$$T_{ewh,\tau}^{up} - \left(T_{ewh}^{set} + 1°\mathrm{C}\right) \leq \; z_\tau^{comf}\;, \tag{24}$$

are introduced to penalise deviations from consumer set-points for the EWH water T_{ewh}^{set}.

5.4. Modelling an electric resistance space heater

To model the dynamics of the space heating system, a data-driven model proposed in [34] is adopted. Similarly, we replace the coefficient of performance (COP) with one, to match the resistance-based ERH. The model coefficients are derived by using statistical learning on data recorded from actual heating systems. The following constraint captures the dynamics of the system:

$$T_{air,\tau+1}^{in} = T_{air,\tau}^{in} + \gamma_1\left(T_{air,\tau}^{out} - T_{air,\tau}^{in}\right) + \gamma_2\left(U_{erh,\tau} P_{erh}^{\max}\right) + \gamma_3 P_{irr,\tau} \tag{25}$$

where γ_1, γ_2, and γ_3 are parameters learned from data, $T_{air,\tau}^{in}$ and $T_{air,\tau}^{out}$ are the indoor and outdoor temperatures at time τ, respectively, $U_{erh,\tau} \in [0,1]$ is the ERH duty cycle, $P_{irr,\tau}$ is the solar irradiance at time τ, and $P_{erh}^{\max}$ is the rated power of the ERH. Similar to the EWH, the proxy comfort variables z_τ^{comf} are used to penalise deviations from

consumer set-points. However, as deviations in indoor temperature affect consumer comfort to a higher degree than hot water temperatures, deviations (per °C) are penalised with a larger coefficient:

$$10\left[\left(T^{set}_{erh} - 1°\mathrm{C}\right) - T^{in}_{air,\tau}\right] \leq z^{comf}_{\tau} \tag{26}$$

$$10\left[T^{in}_{air,\tau} - \left(T^{set}_{erh} + 1°\mathrm{C}\right)\right] \leq z^{comf}_{\tau}\ , \tag{27}$$

where T^{set}_{erh} is the consumer indoor temperature set-point.

5.5. Optimisation problem for an ESS-based HEMS controller

For an ESS-based HEMS controller, the following objective function is used:

$$\begin{aligned} &\underset{y_\tau, z^{ij}_\tau}{\text{minimise}} \quad \frac{1}{W+1}\sum_{\tau=w}^{w+W} c_\tau y_\tau + \mu_w \tilde{I}(X_w; Y_w) \\ &\text{subject to} \quad (y_\tau,\ z^{ij}_\tau) \in \mathcal{F}_{ess,\tau}\ , \end{aligned} \tag{28}$$

where c_τ is the cost of energy, μ_w is the price-of-privacy-loss, and the set $\mathcal{F}_{ess,\tau}$ enforces constraints (1), and (9) to (16).

The inclusion of the energy costs penalises the charging of the ESS during high-price periods, and when coupled with lower prices-of-privacy-loss, discourages multiple charge–discharge cycles within a day. This allows a better comparison with FTL-based systems, which cannot 'discharge', and hence have equivalent energy storage capacities limited by the average daily thermal demand and system losses.

5.6. Optimisation problems for FTL-based HEMS controllers

In addition to the energy costs, the optimisation objective for FTLs should also minimise consumer comfort violations. We minimise $\|\mathbf{z}^{comf}\|_2^2$, $\mathbf{z}^{comf} := [z^{comf}_w, z^{comf}_{w+1}, \ldots, z^{comf}_{w+W}]^\top$, which imposes larger penalties for larger comfort violations. Thus, the optimisation problem for an EWH-based HEMS controller is given by

$$\begin{aligned} &\underset{y_\tau, z^{ij}_\tau, \mathbf{z}^{comf}}{\text{minimise}} \quad \frac{1}{W+1}\sum_{\tau=w}^{w+W} c_\tau y_\tau + \mu_w \tilde{I}(X_w; Y_w) + \rho_w \|\mathbf{z}^{comf}\|_2^2 \\ &\text{subject to} \quad (y_\tau,\ z^{ij}_\tau, \mathbf{z}^{comf}) \in \mathcal{F}_{th,\tau}\ , \end{aligned} \tag{29}$$

where ρ_w is the consumer comfort coefficient, and the set $\mathcal{F}_{th,\tau}$ enforces the constraints (1), (9) to (11), and (17) to (24). For a system with both an EWH and an ERH, set $\mathcal{F}_{th,\tau}$ in (29) is replaced with the set $\mathcal{F}'_{th,\tau}$, which now also includes constraints (25) to (27).

6. Numerical experiments

House 23618 from the Residential Building Stock Assessment (RBSA) database [35] was arbitrarily chosen and used for the numerical simulations. This house is based in Emmett, Idaho, USA, which has a semi-arid climate with cold winters and multiple heating-days. Weather data with 5-minute resolution from Boulder, Colorado, USA, which has a similar climate, was used in the simulations. The HEMS controllers from Section 5 were simulated for 180 heating-days with hourly resolution in MATLAB 2018a and the Gurobi 8.1.0 optimisation solver.

For simplicity, we assume that the incoming water supply temperature is constant, and that μ_w and ρ_w, which can be time-dependent, are also constant. Moreover, for ease of comparison, we assume that the controller has perfect knowledge of the sensitive load across the prediction horizon, and that the models used in the controller accurately represent the actual systems. The equivalent energy storage capacity of an FTL is hard to estimate, depends on many stochastic parameters such as weather conditions and consumer behaviour, and remains an ongoing research challenge. For the simulations, we assumed that this capacity is given by the average daily thermal demand of the household over the simulation period, considering the simulation setup and assumptions. The general simulation parameters are given in Table 1,

Table 1
General parameters used in the simulations.

Prediction horizon, $W+1$	24
MI approximate sample size,[a] N_ε	201.6
Number of $\mathcal{X}$ Bins, m	24
Number of $\mathcal{Y}$ Bins, n	24
Energy price (peak)	24.6 cents/kWh
Energy price (off-peak)	13.15 cents/kWh
Minimum grid load, $Y^{\min}$	0 kW
Maximum grid load, $Y^{\max}$	12 kW

[a]Including additive smoothing constant.

Table 2
System-specific simulation parameters.

	ESS	EWH	ERH
Equivalent storage cap.	6.29 kWh	6.29 kWh	32.63 kWh
Power rating	5.5 kW	5.5 kW	4.5 kW
1-way efficiency/COP	96%	1	1
Absolute min. temp.	–	50 °C	–
Absolute max. temp.	–	90 °C	–
Consumer set-point	–	75 °C	22 °C
Mains water temp.	–	10 °C	–
Water heat cap., C_p	–	4.19 kJ/K	–
C^{low}_{ewh}	–	356.15 kJ/K	–
C^{up}_{ewh}	–	356.15 kJ/K	–
Thermal coeff., UA^{low}_{ewh}	–	5.82e−4 kW/K	–
Thermal coeff., UA^{up}_{ewh}	–	5.82e−4 kW/K	–
γ_1	–	–	1.50e−2
γ_2	–	–	1.86e−1
γ_3	–	–	3.45e−1

while Table 2 gives the system specific parameters. For the FTL-based controllers, $\rho_w = 10$.

The majority of EWHs and ERHs that are currently installed are step loads, while the thermodynamics of the systems are, in reality, continuous. Smaller simulation step sizes would better capture the actual system dynamics, at the expense of computational tractability. Hence, to better match realistic systems and illustrate the mismatch between optimisation models and reality, the continuous duty-cycles from the hourly HEMS controllers were also converted into 5-minute on–off cycles by a secondary controller for system dynamics simulations. This controller attempts to match the HEMS' duty cycle, whilst also enforcing the FTL constraints in Section 5 at 5-minute resolution. Note that the accuracy of the thermodynamic models is beyond the scope of this paper. To further explore the privacy-protection of both ESS- and FTL-based systems, HEMS controllers that do not consider energy costs were also simulated.

Fig. 2 shows the load profiles from an ESS-based system and an EWH-based system with discretised control actions (5-minute simulation interval), with $\mu_w = 5$, and considering energy costs. As illustrated, the reduced flexibility of the EWH-based system limits its ability to mask the sensitive load, resulting in more instances where the sensitive load is revealed, e.g., around time steps 2866, 2893 and 2916 (highlighted in grey). The ESS is also shown to have a single charge–discharge cycle within 24 h. Note that here, $H(X;Y)$ is maximised instead as it was impossible to achieve minimal $H(Y)$.

Quantitatively, the privacy leakage of the various systems were assessed by first treating the loads as i.i.d. (IID MI) processes, and then as stationary first-order Markov processes (Markov MI), using the MI estimation methods described in [31]. It is important to note that the MI estimation methods assume that the FTLs are not privacy-sensitive, *i.e.*, the privacy leakage from the FTL use is not considered. This is particularly important when interpreting the results for $\mu_w = 0$ and energy cost is not considered in the objective function. Table 3 summarises the MI estimates from the various systems.

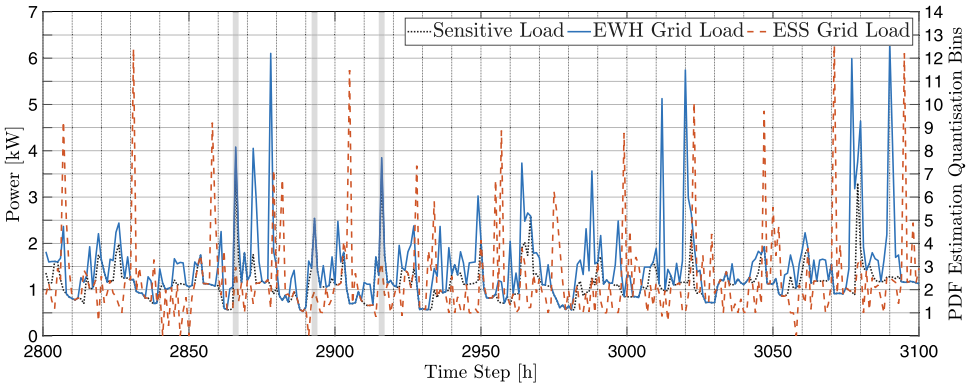

Fig. 2. The sensitive load, and grid loads with $\mu_w = 5$, illustrating how the privacy protection algorithm has shaped the grid visible load to mask the information in the sensitive load, e.g., load peaks and troughs.

Both the ESS- and EWH-based systems reached their privacy protection limit without sacrificing the other objectives with $\mu_w = 5$. As seen, the ESS system has less than half the privacy leakage compared to the EWH system with $\mu_w > 0$. Without considering energy costs, it can be seen that the ESS achieves much lower MI values (with multiple charge–discharge cycles within a day), while there is only marginal improvement for the EWH due to comfort considerations. With a maximum water draw of 112 litres within an hour from the 170 litre hot water tank, there is insufficient flexibility when using the EWH to protect privacy with a 1°C dead-band. Due to the limited operational flexibility afforded by the hot water tank size and safety considerations, the operation of the EWH does not vary by much given different values of μ_w. This can be seen in Fig. 3, which plots the load curves for different values of μ_w, without considering energy cost.

The marginal increase in MI for the ESS without energy costs is due to the binary variables (multiple solution candidates). Moreover, the effects of model mismatch is briefly studied by comparing the 5-minute step load versus non-step load hourly EWH simulations. The actual operation of the EWH differs from the solution of the hourly control actions, as the system dynamics require the secondary controller to make minor adjustments in order to prevent constraint violations (e.g., more accurate water mixing and loss modelling). While one could use models that better represent the continuous dynamics of the thermal system, model mismatch is inevitable in reality, but that is beyond the scope of this paper. The minor adjustments by the secondary controller eventually led to minor reduction in MI in most cases, but that is coincidental.

Even when combining the EWH with an ERH, the privacy protection afforded still falls below that of the ESS with a fraction of the storage capacity for $\mu_w > 0$. More importantly, the use of the ERH for privacy entails a significant Markov MI increase, due to the time-correlated dynamics of the system. While there is substantial MI reduction for all systems even with $\mu_w = 0$ (the i.i.d. entropy/MI for the sensitive load is 2.710 bits), if the EWH and ERH usage is privacy-sensitive, then at $\mu_w = 0$, the EWH and ERH profiles are unprotected and fully reveal the information contained by their usage.

The limitation of the ERH in providing more privacy protection even when energy costs are ignored, again, lies in the fact that the temperature dead-band is 1°C, limiting flexibility. This dead-band prevents over-heating the space or letting it cool below comfortable levels. Note that there is very low IID MI when $\mu_w = 0$ for the combined EWH and ERH system. This is due to the fact that coincidentally, the period when there is high space heating demand is also the period with high private information leakage (occupied and low-load night periods); and that the ERH usage is assumed to not reveal private information. Fig. 4 illustrates the sensitive load and grid load curves of the system with an ERH and EWH, without considering energy cost. As shown, while the FTL system is running most of the time, the peak FTL energy demand coincides with the sensitive load troughs, e.g., between time step 2850 and 2865. Moreover, given comfort considerations, the controller has limited flexibility in rescheduling the FTL energy demand; as reflected by the cumulative energy consumption of the $\mu_w = 10$ curve closely matching that of the comfort-only curve within short periods of time.

In order to get a sense of how an oversized FTL system with more operational flexibility could affect privacy protection, the EWH-based system was simulated with a tank size of 255 litres instead of 170 litres (50% larger). At $\mu_w = 10$ and without considering energy costs, the larger step load EWH-based system achieved slightly better privacy, with an IID MI value of 0.614 bits (versus 0.633 bits). This is true, even when considering energy costs with $\mu_w = 10$ (0.636 versus 0.647 bits). Nonetheless, this slight improvement in privacy protection would probably not justify the over-sizing of the FTL systems (compared to an investment in an ESS for energy cost optimisation and privacy protection).

As shown in the studies conducted in [17], all else being equal, the characteristics of the underlying sensitive load profile affects the level of perceived privacy regardless of privacy metric. To better generalise the findings from the numerical study, the simulations were repeated with the same system parameters for the EWH, ERH and ESS as that used for House 23618, but using the sensitive load profile of House 21355 from the RBSA database. The privacy loss for the House 21355 sensitive load profile, which has an i.i.d. entropy value of 2.246 bits, is shown in Table 4. As can be seen, better privacy protection is achieved for the House 21355 profile. This is due to the various systems, *i.e.*, EWH, ERH and ESS, being better able to match House 21355's lower peak consumption of 4.87 kW (versus 5.22 kW for House 23618), and its lower daily energy consumption. While there is less noticeable improvement for the EWH and ERH-based system relative to one based on an ESS, the findings are similar to those from using the profile of House 23618. This indicates that the findings from the numerical simulations, which validate the conclusion from the previous theoretical analysis in Section 4, are not specific to a particular sensitive load profile.

Table 5 summarises the percentage change in average daily energy cost of the various systems for House 23618 relative to a no-privacy, non-cost-optimised solution. For the FTL-based systems, the cost basis is taken as the "consumer comfort-only" setting, while the energy cost of the original sensitive load is used as the basis for the ESS-based systems. These are highlighted in yellow in Table 5. As can be seen, the controller is able to reduce energy costs given a two-tier price tariff, while simultaneously protecting consumer privacy. Nonetheless, the cost savings are minor for the EWH-based systems due to the operational inflexibility of the EWH. For systems that utilise both the EWH and ERH, the cost savings are more significant (both in terms of absolute and relative amounts). This is due to the higher energy demand of the ERH coupled with the higher thermal inertia in the space heating system. For ESS-based systems, the relative cost savings

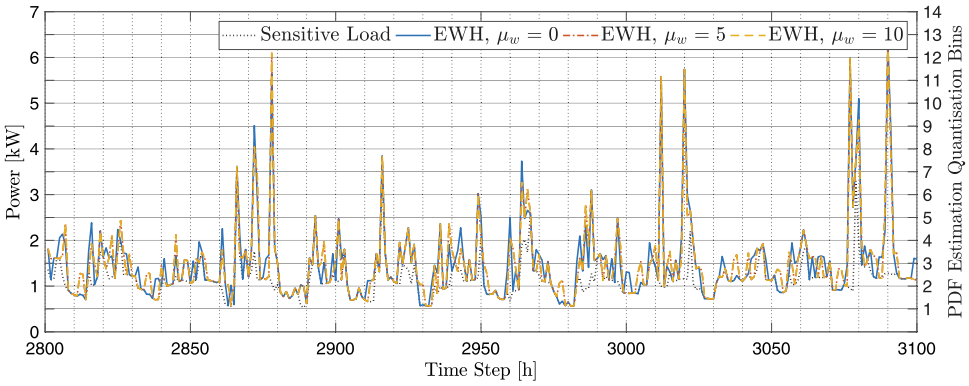

Fig. 3. The sensitive load, and grid loads for EWH schemes without energy cost and with different μ_w values, showing that there is a lack of flexibility in the EWH load due to consumer comfort and safety constraints.

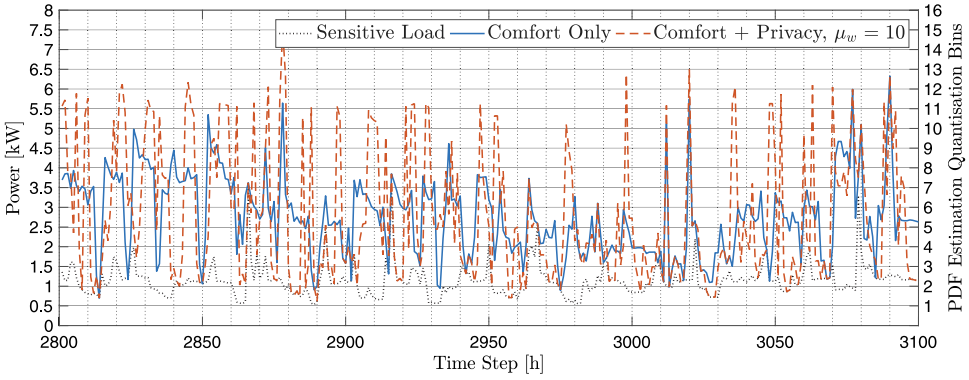

Fig. 4. The sensitive load, and grid loads for ERH + EWH scheme without energy cost. The load profiles show that the energy consumption of the system when protecting consumer privacy matches that of the "comfort-only" setting within a small time window. This illustrates the operational inflexibility of the ERH + EWH system, given the tight temperature dead-bands.

Table 3
Privacy loss of house 23 618 with a 24-hour prediction horizon; illustrating the level of privacy protection afforded by the different systems under different prices-of-privacy-loss.

	$\mu_w = 0$		$\mu_w = 5$		$\mu_w = 10$	
	IID MI	Markov MI	IID MI	Markov MI	IID MI	Markov MI
ESS with energy costs	0.565	0.709	0.286	0.678	0.287	0.653
ESS without energy costs	–	–	0.149	0.672	0.154	0.671
Step load EWH with energy costs	0.656	0.859	0.655	0.837	0.647	0.825
Step load EWH without energy costs	0.657	0.831	0.633	0.817	0.633	0.817
Non-step load EWH with energy costs	0.791	0.941	0.693	0.870	0.679	0.864
Non-step load EWH without energy costs	0.813	0.915	0.628	0.821	0.628	0.824
Step load EWH and ERH with energy costs	0.367	1.062	0.362	1.066	0.362	1.080
Step load EWH and ERH without energy costs	0.136	0.788	0.326	1.214	0.326	1.208

Table 4
Privacy loss of house 21 355 using system parameters from house 23 618 and a 24-hour prediction horizon; illustrating the level of privacy protection afforded by the different systems under different prices-of-privacy-loss.

	$\mu_w = 0$		$\mu_w = 5$		$\mu_w = 10$	
	IID MI	Markov MI	IID MI	Markov MI	IID MI	Markov MI
ESS with energy costs	0.287	0.515	0.223	0.555	0.232	0.539
ESS without energy costs	–	–	0.140	0.564	0.140	0.523
Step load EWH with energy costs	0.530	0.771	0.517	0.766	0.519	0.763
Step load EWH without energy costs	0.528	0.753	0.496	0.738	0.496	0.738
Non-step load EWH with energy costs	0.686	0.876	0.583	0.807	0.563	0.802
Non-step load EWH without energy costs	0.710	0.844	0.513	0.792	0.514	0.792
Step load EWH and ERH with energy costs	0.269	0.989	0.266	0.988	0.266	0.997
Step load EWH and ERH without energy costs	0.120	0.778	0.271	1.070	0.264	1.094

are larger, but the absolute value is still overshadowed by the system's high investment costs. Based on estimates in [36], which are inline with projections in [37], the latest prices for behind-the-meter battery packs (without installation and balance-of-system costs) are between $400 and $750 per kilowatt-hour. This translates to a simple payback period of between 11 and 20 years, which is not attractive given the expected lifetime of batteries. On the other hand, prioritising privacy protection

Table 5
Average daily energy cost of house 23 618 with a 24-hour prediction horizon; showing how the costs change when different systems and prices-of-privacy-loss are used.

	$\mu_w = 0$	$\mu_w = 5$	$\mu_w = 10$
ESS with energy costs	−14.96%	−14.98%	−15.00%
ESS without energy costs[a]	$4.1693	+3.280%	+4.069%
Step load EWH with energy costs	−0.539%	−0.531%	−0.538%
Step load EWH without energy costs	$5.970	+0.270%	+0.270%
Non-step load EWH with energy costs	−1.081%	−1.068%	−1.055%
Non-step load EWH without energy costs	$5.957	+0.102%	+0.102%
Step load EWH and ERH with energy costs	−9.086%	−9.137%	−9.116%
Step load EWH and ERH without energy costs	$12.202	+1.903%	+2.237%

[a]The energy cost of the underlying sensitive load is used as the basis for comparing the ESS-based schemes.

(ignoring energy costs) leads to a slight increase in energy cost, which is comparable across both ESS- and FTL-based systems.

Further experiments are required in order to fully generalise the empirical findings across different load profile characteristics, climate zones, and building and system parameters, which is left as the subject of future work.

7. Conclusions and future outlook

The topic of smart meter consumer privacy is an important one, given that advanced metering infrastructures are often touted as the bedrock of the smart electric grid. Without viable solutions to protect consumer privacy, the deployment of smart meters could potentially be jeopardised. In this paper, we studied the use of resistive flexible thermal-based consumer loads for consumer privacy protection using user demand shaping methods, comparing them against systems using energy storage systems. By conducting a theoretical analysis using mutual information as the quantitative measure of privacy, we show that, based on the fact that flexible thermal-based consumer loads are unable to compensate sensitive load by 'discharging', the level of protection afforded by them is below that of energy storage systems. Moreover, as seen from the numerical experiments, the inflexibility of these systems due to the time-specific nature of thermal demand limits their performance; unless one allows for large temperature fluctuations or use largely over-sized systems. Nonetheless, they are still able to afford some level of privacy protection, with incremental energy costs that are comparable to those of energy storage systems, but without the added upfront investment costs. Coupled with their increasing ubiquity in consumer households, research on utilising flexible thermal-based consumer loads in privacy protection schemes, whether standalone or in conjunction with energy storage systems, should be expanded.

Future work will consider the use of inductive loads and loads with interruptible, but fixed cycle lengths for consumer privacy protection, and the generalisation of the findings using empirical studies.

CRediT authorship contribution statement

Jun-Xing Chin: Conceptualization, Methodology, Formal analysis, Investigation, Software, Visualisation, Writing - original draft. **Kyri Baker:** Conceptualization, Methodology, Formal analysis, Writing - review & editing. **Gabriela Hug:** Conceptualization, Supervision, Writing - review & editing, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] McKenna E, Richardson I, Thomson M. Smart meter data: Balancing consumer privacy concerns with legitimate applications. Energy Policy 2012;41:807–14.
[2] McDaniel P, McLaughlin S. Security and privacy challenges in the smart grid. IEEE Secur Priv 2009;7(3).
[3] Molina-Markham A, Shenoy P, Fu K, Cecchet E, Irwin D. Private memoirs of a smart meter. In: Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building. Zurich, Switzerland: 2010, p. 61–66.
[4] PWC. Consumer intelligence series: Protect.me. Tech. rep., Pricewaterhouse-Coopers; 2017, [online]. Available: https://www.pwc.com/us/en/advisory-services/publications/consumer-intelligence-series/protect-me/cis-protect-me-findings.pdf.
[5] European Union. Regulation (EU) 2016/679 of the European parliament and of the council of 27 april 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46/EC (general data protection regulation). Off J Eur Union 2016.
[6] McLaughlin S, Podkuiko D, Miadzvezhanka S, Delozier A, McDaniel P. Multi-vendor penetration testing in the advanced metering infrastructure. In: Proceedings of the 26th Annual Computer Security Applications Conference. Austin, Texas, USA: ACM; 2010, p. 107–16.
[7] Winter JS. Citizen perspectives on the customization/privacy paradox related to smart meter implementation. Int J Technoethics (IJT) 2015;6(1):45–59.
[8] Van Aubel P, Poll E. Smart metering in the netherlands: What, how, and why. Int J Electr Power Energy Syst 2019;109:719–25.
[9] Sankar L, Rajagopalan SR, Mohajer S, Poor HV. Smart meter privacy: A theoretical framework. IEEE Trans Smart Grid 2013;4(2):837–46.
[10] McLaughlin S, McDaniel P, Aiello W. Protecting consumer privacy from electric load monitoring. In: Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS '11). Chicago, Illinois, USA: 2011, p. 87–98.
[11] Yang W, Li N, Qi Y, Qardaji W, McLaughlin S, McDaniel P. Minimizing private data disclosures in the smart grid. In: Proceedings of the 19th ACM Conference on Computer and Communications Security (CCS '12). Raleigh, North Carolina, USA: 2012, p. 415.
[12] Tan O, Gomez-Vilardebo J, Gündüz D. Privacy-cost trade-offs in demand-side management with storage. IEEE Trans Inf Forensics Secur 2017;12(6):1458–69.
[13] Chin JX, Tinoco De Rubira T, Hug G. Privacy-protecting energy management unit through model-distribution predictive control. IEEE Trans Smart Grid 2017;8(6):3084–93.
[14] Zhang Z, Qin Z, Zhu L, Weng J, Ren K. Cost-friendly differential privacy for smart meters: Exploiting the dual roles of the noise. IEEE Trans Smart Grid 2017;8(2):619–26.
[15] Baker K, Garifi K. Power signature obfuscation using flexible building loads. In: International Workshop on Non-Intrusive Load Monitoring (NILM). Austin, TX, USA: 2018.
[16] Mashima D, Serikova A, Cheng Y, Chen B. Towards quantitative evaluation of privacy protection schemes for electricity usage data sharing. ICT Express 2018;4(1):35–41.
[17] Arzamasov V, Schwerdt R, Karrari S, Böhm K, Nguyen TB. Privacy measures and storage technologies for battery-based load hiding - an overview and experimental study. In: Proceedings of the Eleventh ACM International Conference on Future Energy Systems. e-Energy '20, Australia: Virtual Event; 2020, p. 178–95.
[18] Giaconi G, Gündüz D, Poor HV. Privacy-aware smart metering: Progress and challenges. IEEE Signal Process Mag 2018;35(6):59–78.
[19] Danezis G, Fournet C, Kohlweiss M, Zanella-Béguelin S. Smart meter aggregation via secret-sharing. In: First ACM Workshop on Smart Energy Grid Security SEGS '13. Berlin, Germany: 2013.
[20] Koo D, Shin Y, Hur J. Privacy-preserving aggregation and authentication of multi-source smart meters in a smart grid system. Appl Sci 2017;7(10).
[21] Mustafa MA, Cleemput S, Aly A, Abidin A. A secure and privacy-preserving protocol for smart metering operational data collection. IEEE Trans Smart Grid 2019;10(6):6481–90.
[22] Efthymiou C, Kalogridis G. Smart grid privacy via anonymization of smart metering data. In: 2010 First IEEE International Conference on Smart Grid Communications. Gaithersburg, MD, USA: 2010.
[23] Rottondi C, Mauri G, Verticale G. A data pseudonymization protocol for smart grids. In: 2012 IEEE Online Conference on Green Communications (GreenCom). 2012.
[24] Ni J, Zhang K, Alharbi K, Lin X, Zhang N, Shen XS. Differentially private smart metering with fault tolerance and range-based filtering. IEEE Trans Smart Grid 2017;8(5):2483–93.
[25] Acs G, Castelluccia C. I have a DREAM! (DiffeRentially PrivatE smart Metering). In: The 13th Information Hiding Conference (IH). 2011, p. 118–32.
[26] Hassan MU, Rehmani MH, Kotagiri R, Zhang J, Chen J. Differential privacy for renewable energy resources based smart metering. J Parallel Distrib Comput 2019;131.
[27] Giaconi G, Gündüz D, Poor HV. Smart meter privacy with renewable energy and an energy storage device. IEEE Trans Inf Forensics Secur 2018;13(1):129–42.

[28] Chen D, Kalra S, Irwin D, Shenoy P, Albrecht J. Preventing occupancy detection from smart meters. IEEE Trans Smart Grid 2015;6(5):2426–34.
[29] Sun Y, Lampe L, Wong VWS. Smart meter privacy: Exploiting the potential of household energy storage units. IEEE Internet Things J 2018;5(1):69–78.
[30] Zhou B, Li W, Chan KW, Cao Y, Kuang Y, Liu X, Wang X. Smart home energy management systems: Concept, configurations, and scheduling strategies. Renew Sustain Energy Rev 2016;61:30–40.
[31] Chin JX, Giaconi G, Tinoco De Rubira T, Hug G, Gündüz D. Considering time correlation in the estimation of privacy loss for consumers with smart meters. In: 2018 Power System Computation Conference (PSCC). Dublin, Ireland: 2018.
[32] McLoughlin F, Duffy A, Conlon M. The generation of domestic electricity load profiles through Markov chain modelling. Euro-Asian J Sustain Energy Dev Policy 2010;3(January - December).
[33] Jin X, Maguire J, Christensen D. Model predictive control of heat pump water heaters for energy efficiency. In: 2014 ACEEE Summer Study on Energy Efficiency in Buildings. Pacific Grove, CA: 2014.
[34] Jin X, Baker K, Christensen D, Isley S. Foresee: A user-centric home energy management system for energy efficiency and demand response. Appl Energy 2017;205:1583–95.
[35] Ecotope Inc. Residential building stock assessment: Metering study. Tech. rep. E14-283, 2014, [Online]. Available: https://neea.org/img/documents/residential-building-stock-assessment-metering-study.pdf.
[36] EnergySage LLC. How much does solar storage cost? Understanding solar battery prices. 2020, https://www.energysage.com/solar/solar-energy-storage/what-do-solar-batteries-cost/ (accessed 19 July 2020).
[37] IRENA. Innovation landscape brief: Behind-the-meter batteries. Tech. rep., International Renewable Energy Agency; 2019, [Online]. Available: https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2019/Sep/IRENA_BTM_Batteries_2019.pdf.