A Mechanism for Cooperative Demand-Side Management

Guangchao Yuan*, Chung-Wei Hang†, Michael N. Huhns‡ and Munindar P. Singh*
*North Carolina State University, Raleigh, NC, USA
email: {gyuan.ncsu@gmail.com, m.singh@ieee.org}
†IBM; email: chungwei.hang@gmail.com
‡University of South Carolina; email: huhns@sc.edu

***Abstract*—Demand-side management (DSM) is an important theme in studies of the Smart Grid and offers the possibility of leveling power consumption with its attendant benefits of reducing capital expenses. This paper develops an algorithmic mechanism that reduces peak total power consumption and encourages *prosocial* behavior, such as expressing flexibility in one's power consumption and reporting preferences truthfully. The objective is to provide a tractable, budget-balanced mechanism that promotes truth-telling from households. The resulting mechanism is theoretically and empirically proven to be *ex ante* budget-balanced, weakly Pareto-efficient, and weakly Bayesian incentive-compatible. A simulation study verifies that the mechanism could largely reduce the computational complexity that the optimal allocation requires, while maintaining approximately the same performance. A user study with 20 subjects further shows the effectiveness of the mechanism in preventing participants from defecting and incentivizing them to reveal flexible preferences.**

***Index Terms*—Demand-side management; Prosocial behaviors**

I. Introduction

Demand-side management (DSM), which is based on the idea of controlling demand to match supply, is a key aspect of the Smart Grid. The objective of DSM is to reduce peak power consumption by shifting peak-hour demand to off-peak hours. Current DSM approaches involve either turning household appliances on or off directly via a smart meter or applying real-time pricing [1]. However, directly using a load control signal to turn off appliances can lead to consumer dissatisfaction. Real-time pricing is difficult for consumers [2] and can lead to chaotic outcomes [3].

With advances in computing power and reliable two-way communication infrastructures, it becomes feasible to employ autonomous agents to capture preferences of both resource providers and consumers, and represent the interactions between them. Therefore, instead of focusing on how each consumer interacts with a resource provider individually, researchers have begun to study approaches with aggregate load management from multiple consumers. Existing approaches fall into two main categories. Those in the first category assume that consumers' objectives are aligned with resource providers' objectives and model the problem as a global optimization problem [4], [5]. Those in the second category model the problem in a game theory setting where selfish consumers compete with each other to minimize their costs [6], [7].

However, previous studies suffer from two main limitations. First, most existing approaches either assume that consumers' preferences are revealed truthfully [8] or require a considerable amount of computation [9]. On the one hand, consumers may be unwilling to reveal such private information. On the other hand, the computational complexity of an approach would preclude large-scale systems.

Second, most prior works do not consider whether the budget of resource providers could remain balanced. *Budget-balance* is an important property that guarantees a resource provider won't run into deficit; i.e, the resource provider's purchase payment does not exceed the sum of payments from all of its consumers. At the macro-level, a wholesale power market [10] functions as a single-sided auction where resource providers bid for a given amount of power for the next day and wholesale prices are lower during off-peak periods. At the micro-level, each consumer pays charges to one or more resource providers for its consumption, and each consumer pays a smaller bill when peak consumption is lowered. Failure to ensure the budget-balance of a resource provider may reduce the practicality of the system.

The well-known Vickrey-Clarke-Groves or VCG mechanism [11] can address the first limitation, but not the second. VCG is a pricing method that motivates each rational household (or an agent, in general) to reveal its true valuation. Each household is charged according to the true social cost it imposes on others and thus has an incentive to declare its valuation truthfully. However, resource providers may run a deficit, which is not acceptable in settings such as residential DSM. In addition, a VCG mechanism may not be tractable—under the usual assumption that a mechanism is *tractable* when the computation it requires can be carried out in polynomial time [11].

Inspired by the VCG mechanism, we propose and investigate an algorithmic mechanism in a Bayesian game setting [12], dubbed *Enki*. Enki takes advantage of consumers' flexibility [13] in consuming power (a flexible consumer provides more choices to Enki to schedule his or her consumption) and encourages such prosocial behavior (flexible consumers pay less), through Enki's allocation objective and payment mechanism. In this way, Enki aligns the interest of each

Guangchao Yuan is now with Microsoft.

household with the overall goal of lowering peak total consumption. Enki is a tractable mechanism guaranteeing that resource providers would not run into a deficit. Further, Enki still provides incentives to households to act according to their true valuations, except that we relax the truth-telling from dominant strategies in VCG to Bayes-Nash equilibrium strategies [14].

Enki is a novel DSM method for the day-ahead energy market setting. The method applies at a collective, i.e., neighborhood, level. The neighborhood functions as a resource provider. Every day each household (consumer) declares its expected consumption for the next day, specifying the amount to be consumed and a time period within which it would be willing to confine its consumption. The report can be made through an Energy Consumption Controller (ECC) unit embedded in each household's smart meter. An ECC unit functions as a decision-making tool. Specifically, it operates as follows:

- learn each household's daily power consumption pattern through machine learning techniques;
- decide; and
- report the household's demand for the next day.

An ECC unit connects with a neighborhood controller through a local network [6], [9]. The neighborhood allocates time ranges to each household within which it can consume power. The neighborhood ensures the allocation satisfies *all* consumers' requests while reducing peak demand. Each household is charged with a bill for its actual (i.e., not merely predicted) consumption.

Contributions

Enki's novelty lies in, first, how it exploits consumers' flexibility in power consumption to align their interests with the group's interest. Second, Enki is a tractable mechanism guaranteeing that resource providers (neighborhoods) would not run into a deficit; Enki is weakly Bayesian incentive-compatible and weakly Pareto-efficient. Third, we conduct a simulation study verifying that the greedy allocation in Enki approximates the optimal allocation while reducing computation time. Evaluation results from a user study show that Enki is effective in incentivizing participants to tell the truth and be flexible given their preferences.

II. Background

Demand-side management (DSM), which seeks to reduce peak demand, is a longstanding theme in the power industry. Current DSM approaches are of two types [15]. One, Direct Load Control (DLC) involves a power company turning off selected appliances (e.g., HVAC) during peak hours. Consumers often find ceding such control to a power company risky since their particular needs might not be addressed. And, participation in such programs appears low, at least in anecdotal terms [16]. Two, Price-Based Control (PBC) involves a power company incentivizing a consumer to deter peak-time electricity usage [17]. However, PBC has the drawback of often shifting the peak from one period to another [18]. Because consumers often respond to a price signal, they all tend to shift to the lowest price period without a controller.

Researchers use agents that take into account consumers' preferences to assist them in making decisions under different pricing models. One category of existing approaches involves an optimal allocation by maximizing social welfare [19], [20], [21], [22], [5]. Guo et al. [4] propose a decentralized algorithm to minimize the total energy cost within a neighborhood, where each household may have different energy loads. However, they do not consider consumers' preferences.

A second category of approaches apply game theory. They assume each consumer is self-interested and competes with others to minimize costs. Mohsenian-Rad et al. [6] present a game-theoretic framework to analyze the energy consumption scheduling problem by assuming that the energy requirement of each consumer is determined in advance. Ramchurn et al. [23] propose a decentralized adaptive mechanism to defer each consumers' power usage based on their preference profile. In this way, consumers coordinate to reduce the overall peak consumption. Other works using an adaptive mechanism employ game-theoretic analyses and assume that each consumer is rational [8]. Chen et al. [7] formulate an aggregate game to model the strategic behaviors of selfish consumers. The above studies focus on changing consumer behaviors. They assume that each consumer reveals his or her preferences truthfully. However, consumers may be unwilling to reveal such private information or not know how to reveal it.

Rose et al. [24] outline a mechanism to elicit true consumer preferences for the wholesale market. A neighborhood purchases power from the market using each consumer's prediction and allocates power to each consumer based on the consumer's real need. The neighborhood is charged for any imbalance between the amount it purchased and the aggregate amount that the neighborhood's consumers consumed. Rose et al.'s mechanism is weakly budget-balanced and incentive-compatible. Stein et al. [25] establish similar results for electric vehicle charging. Some researchers propose an incentive mechanism to elicit truthful demand response from geo-distributed data centers [26], [27].

Nevertheless, residential DSM to reveal true preferences from consumers has garnered little attention. Notably, Samadi et al. [9] propose a VCG mechanism for DSM and show it satisfies the properties of *efficiency*, *user truthfulness*, and *nonnegative transfer*. As for other VCG mechanisms, the computational complexity can be NP-hard in many practical domains [11], and Samadi et al.'s mechanism does not scale to a large community.

Instead of using the VCG mechanism, Enki takes advantage of households' flexibility and designs a tractable, budget-balanced mechanism that aligns each household's self-interest with the neighborhood's overall objective. Further, Enki is evaluated by a study with human subjects—such studies being rare in this literature.

III. The Enki Neighborhood Model

A *neighborhood* agent works with several *household* agents by mediating between them and the *power company*. Figure 1 shows how this model works, and Table I includes important notations.

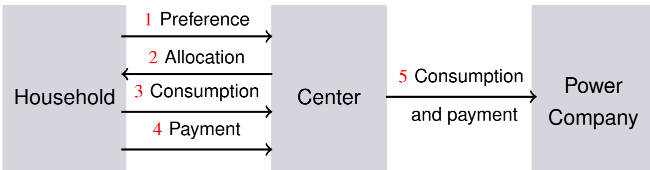

Fig. 1: Neighborhood model architecture.

TABLE I: Notation used in the paper

Symbol	Description
I	the set of households
$H = \{0, \ldots, 23\}$	hours of a day
α_i	true preferred beginning time
β_i	true preferred ending time
$\hat{\alpha}_i$	reported preferred beginning time
$\hat{\beta}_i$	reported preferred ending time
v_i	preferred duration
$\chi_i = (\alpha_i, \beta_i, v_i)$	true preference
$\hat{\chi}_i = (\hat{\alpha}_i, \hat{\beta}_i, v_i)$	reported preference
$s_i = (\alpha_i^s, \beta_i^s)$	suggested allocation
$\omega_i = (\alpha_i^\omega, \beta_i^\omega)$	real consumption
l_h	consumption aggregated over households
σ	a scaling factor in pricing function
$P_h(l_h)$	the cost of the neighborhood at $h \in H$
$\kappa(\boldsymbol{\omega})$	the cost of the neighborhood
d_i	deferment of a household's consumption
r	the power rating in kW
γ_h	signal value (turning on power load or not)
ρ_i	a household's valuation factor
$\theta_i = (\chi_i, \rho_i)$	the type of a household
τ_i	overlap between allocation and true preference
$V_i(\tau_i, v_i, \rho_i)$	the valuation of a household
f_i	flexibility score of a household
o_i	overlap ratio of consumption and allocation
δ_i	defection score of a household
Ψ_i	social-cost score of a household
k	scaling factor in social-cost function
p_i	payment of a household
ξ	scaling factor in payment function
$U_i(\omega_i \vert \chi_i, \theta_i)$	the quasilinear utility of a household
F_i	normalized flexibility score of a household
Δ_i	normalized defection score of a household
b_i	power consumption of a household
p_i^z	household's payment without Enki
$\kappa(\boldsymbol{\omega}^{\boldsymbol{z}})$	cost to the neighborhood without Enki
U_i^z	quasilinear utility of a household without Enki
n	number of households

Preference: Through its ECC, each household reports its predicted consumption for the next day to the neighborhood (center). We abstract the load to a single value. (One can imagine that the households are describing when they will use a notional appliance.) One possible application could be charging electric vehicles. In addition, we focus on *shiftable loads* [23] since the operation of *nonshiftable loads* cannot be scheduled and its cost cannot be reduced. However, our model can be easily extended to a more concrete scenario by considering several such preferences for a given household and adding a constant cost to each household's payment. We simplify our model for clarity in our mathematical presentation and simulation.

Let I represent the set of households. Further, let $H = \{0, 1, \ldots, 23\}$ be the hours of a day. For each household $i \in I$, let $\chi_i = (\alpha_i, \beta_i, v_i)$ and $\hat{\chi}_i = (\hat{\alpha}_i, \hat{\beta}_i, v_i)$ denote its true and reported preferences, respectively, for consuming power for the following day. Because the duration is fixed or varies in a limited range based on the types of appliances in most cases, we assume that all the households truthfully report the duration. α_i ($\hat{\alpha}_i$) $\in H$ is the true (reported) preferred beginning time, β_i ($\hat{\beta}_i$) $\in H$ is the true (reported) preferred ending time, and v_i is the preferred duration. The household's true (reported) preferred interval is defined as $[\alpha_i, \beta_i]$ ($[\hat{\alpha}_i, \hat{\beta}_i]$). Clearly, $\beta_i - \alpha_i \geq v_i$ and $\hat{\beta}_i - \hat{\alpha}_i \geq v_i$. For example, household i reports its current preference as $\hat{\chi}_i = (18, 22, 2)$, which means that it wants to consume power for two hours at any time between 6PM and 10PM. Below, we abbreviate "true preferred interval" as "true interval" and "reported preferred interval" as "reported interval."

Allocation: Given all the reported preferences, the neighborhood creates allocations for the households, so that (1) each household's allocation is scheduled within its *reported interval*, and (2) the peak consumption is reduced. The neighborhood takes $\hat{\boldsymbol{\chi}}$ as input, and outputs allocations $\mathbf{s} = \{s_i \vert i \in I\}$, where $s_i = (\alpha_i^s \in H, \beta_i^s \in H)$ indicates that it suggests household i consume power from α_i^s to β_i^s. Note that $\beta_i^s - \alpha_i^s = v_i$, $\hat{\alpha}_i \leq \alpha_i^s$ and $\beta_i^s \leq \hat{\beta}_i$. We do not change the duration v_i household i prefers.

Household Consumption: Each household may or may not consume power according to the suggested allocation. The consumption of each household i is denoted as: $\omega_i = (\alpha_i^\omega, \beta_i^\omega)$, where $\beta_i^\omega - \alpha_i^\omega = v_i$. If household i consumes power fully according to the allocation, $\omega_i = s_i$. Otherwise, household i defects but within its *true interval*, i.e., overrides the allocation $\omega_i \neq s_i$ ($\alpha_i \leq \alpha_i^\omega$, $\beta_i^\omega \leq \beta_i$). Ideally, if a household has revealed its true preference χ_i, it will not defect, since its allocation is compatible with its preference.

Center and Household Payments: Following Mohsenian-Rad et al. [6], we adopt a superlinear (quadratic) pricing function, $P_h(l_h) = \sigma l_h^2$, where l_h (in kWh) is the aggregated consumption from all the households at $h \in H$ based on $\boldsymbol{\omega}$, and $\sigma > 0$ is a scaling factor. The superlinearity of price provides a basis for improving social welfare by merely reducing peak consumption. The rationale of the pricing function lies in two assumptions: the cost of generating power increases as the aggregate load increases, and it should be strictly convex. Whereas other forms of convex functions meet the two assumptions (e.g., a two-step piecewise function, as suggested in [6]), a quadratic pricing model is more tractable for the purpose of optimization. The price paid by the neighborhood to the power company is defined by $\kappa(\boldsymbol{\omega})$:

$$\kappa(\boldsymbol{\omega}) = \sum_{h \in H} P_h(l_h) = \sum_{h \in H} \sigma l_h^2 \tag{1}$$

Accordingly, each household should make a payment p_i to the neighborhood based on its consumption. How to calculate p_i is discussed in Sections IV-B2 and IV-B3.

IV. The Enki Mechanism

In residential DSM, every consumer prefers to consume power at will (at his or her convenience) and most consumers would like to consume it during peak hours. Schneier [28] terms such a tension between group and self-interest a societal dilemma. If every consumer acts according to its self-interest, a *tragedy of the commons* can result [29]. For example, if all the consumers choose to consume power during peak hours, the supply might not keep up with demand, leading to a power blackout. An intuitive way to address the societal dilemma [28] is to align each household's interest with the group interest. Enki addresses this problem by incentivizing households to behave prosocially in two ways: (1) revealing true preferences and (2) requesting flexible allocations. Specifically, Enki incentivizes households through (1) its allocation objective and (2) its payment mechanism.

A. Optimization and Greedy Allocation

Producing allocations can be formalized as an optimization problem. For each household $i \in I$, let $d_i \in H$ denote the deferment of consumption from its reported preferred beginning time, and r is the power rating in kW. The objective of the allocation is to choose d_i for each $i \in I$ so as to minimize the neighborhood's cost, where

$$\begin{aligned} &\min_{d_i i \in I} \sum_{h=1}^{24} P_h \left(\sum_{i \in I} \gamma_h \cdot r \right) \\ &s.t. \quad \gamma_h = \begin{cases} 1 & \text{if } (\hat{\alpha_i} + d_i) \leq h < (\hat{\alpha_i} + d_i + v_i) \\ 0 & \text{otherwise} \end{cases} \\ &\qquad 0 \leq d_i \leq (\hat{\beta_i} - \hat{\alpha_i} - v_i) \end{aligned} \tag{2}$$

such that $\gamma_h \in \{0, 1\}$ determines whether the power load is scheduled for that time interval. Note that r will vary when we model multiple appliances for a given household.

The quadratic objective function results in high complexity of computing the optimal allocation [23]. Thus, we adopt a greedy allocation (Section IV-C) method: we consider households in an order based on a heuristically calculated flexibility, allocating the time slots for each household in a way that greedily minimizes the peak load for the households handled so far.

B. Utility of a Household

For each household, we define a quasilinear utility function [11], corresponding to a valuation (willingness to pay) minus a payment. A household's valuation depends on its type θ_i: its true preference χ_i and its *valuation factor* ρ_i (a relative measure of its willingness to pay for the same allocation). Therefore, $\theta_i = (\chi_i, \rho_i)$. A household's *reported valuation* is known to Enki after it has declared its type, and a misreported type leads to a misreported valuation.

1) Valuation: We propose a valuation function $V_i(\tau_i, v_i, \rho_i)$ for each household, where $\tau_i \in [0, v_i]$ is the number of time slots the allocation satisfies a household's *true preference*. We identify these criteria for our valuation function [9], [30]:

- The valuation increases with τ_i until τ_i equals v_i, whence the household's valuation is constant.
- The valuation increases with v_i.
- The valuation increases with ρ_i.
- The marginal benefit of τ_i is nonincreasing.

Thus, we adopt this valuation function for household i:

$$V_i(\tau_i, v_i, \rho_i) = -\frac{\rho_i}{2v_i}(\tau_i)^2 + \rho_i \tau_i, \quad \tau_i \in [0, v_i] \tag{3}$$

$V_i(\tau_i, v_i, \rho_i)$ reaches its maximum value $((\rho_i v_i)/2)$ when τ_i is equal to v_i. If a household has reported its true preference, Enki can ensure its true valuation $V_i(\tau_i, v_i, \rho_i)$ is maximized, since Enki's allocation is always scheduled within the household's preferred interval. Note that a household's valuation need not be revealed to Enki explicitly; revealing a preference from a household is equivalent to revealing its valuation.

2) Payment Mechanism: A VCG mechanism [11], [9] would motivate each household to reveal its true valuation and thus pay according to the *social cost* it imposes on others. However, VCG can be intractable. We seek a payment mechanism that approximately captures the above intuition but is less complex. Conceptually, we model a household's social cost in terms of its truthfulness of preference reporting and its flexibility in consuming power.

If a household reports its true preference, its reported preference indicates its social cost. We introduce *flexibility* f_i to model the effect of household i's reported (true) preference $\hat{\chi_i}$ ($\hat{\chi_i} = \chi_i$) in peak reduction. That is, given that all the households fully follow their allocations $\mathbf{s}$, a more flexible reported preference $\hat{\chi_i}$ has a greater contribution in reducing peak consumption. We say household i is more flexible if it provides a more flexible reported (true) preference $\hat{\chi_i}$, i.e., a higher f_i, and, correspondingly, it should pay less.

Conversely, if a household has misreported its true preference, it will defect, i.e., override s_i, as long as the allocation does not satisfy its true preference. In this case, the household's real consumption reflects its social cost. We introduce *defection* to model the effect of a disobedient household in increasing peak consumption. That is, a household i with greater defection δ_i exerts a greater effect in harming the neighborhood and it should pay more.

The payment mechanism should respect the following properties regarding flexibility and defection.

Property 1: All else being equal, a household that truthfully reports a wider preferred interval would pay less.

Property 2: All else being equal, a household that is truthfully more willing to consume power during off-peak hours would pay less.

Property 3: All else being equal, a household that deviates from its allocation pays more than one that does not.

For motivation, consider three households A, B, C.

Example 1: Suppose each household expresses its true preference of $(18, 20, 1)$, which means each prefers to consume power for one hour from 6PM to 8PM. Because their preferences are the same, they should get equal payments.

Example 2: Suppose each household expresses its true preference, but they differ: $\chi_{\mathsf{A}} = (18, 19, 1)$ and $\chi_{\mathsf{B}} = \chi_{\mathsf{C}} = (18, 20, 1)$. A has a higher probability of leading to a higher peak due to its narrower true interval. Therefore, A is less flexible and should pay more than B and C.

Example 3: Suppose each household expresses its true preference, but they differ: $\chi_{\mathsf{A}} = (16, 18, 2)$ and $\chi_{\mathsf{B}} = \chi_{\mathsf{C}} = (18, 21, 2)$. Even though the true interval of B and C is one hour wider, A prefers an off-peak interval. Thus, A contributes more to reducing the peak than the others.

In the allocation step of Section IV-A, the order in which the households are considered is a permutation of A, B, and C. The resulting allocations are shown in Figure 2:

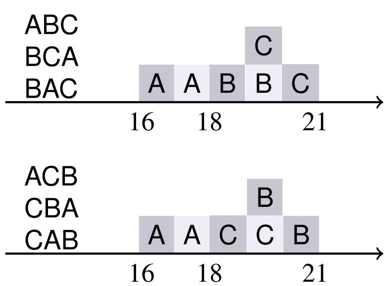

Fig. 2: Allocations corresponding to Example 3.

A's allocation remains the same, and it does not lead to the peak consumption in any permutation. However, B and C have a 50% probability of leading to the peak consumption. Hence, A is more flexible and should pay less.

Example 4: Suppose A and B report the same preference: $\hat{\chi_{\mathsf{A}}} = \hat{\chi_{\mathsf{B}}} = (18, 20, 1)$. In the resulting allocation (Figure 3), A takes the first hour whereas B takes the second. However, B defects whereas A does not. Then, B should obtain a higher defection score and A should pay less.

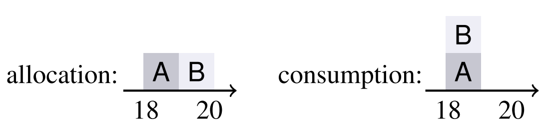

Fig. 3: Household B defects and thus pays more.

3) Measuring Flexibility and Defection: We define the *flexibility* score f_i of household i as

$$f_i = \frac{\beta_i - \alpha_i}{v_i} \frac{1}{N_i}, \tag{4}$$

where $N_i = \frac{\sum_{h=\alpha_i}^{\beta_i} n_h}{\beta_i - \alpha_i}$ is the average number of households (including itself) consuming power within each hour of its preferred interval and n_h is the number of households consuming power in time slot $h \in H$. In Example 2, $N_{\mathsf{B}} = \frac{3+2}{20-18} = 2.5$, and $f_{\mathsf{B}} = \frac{20-18}{1} \frac{1}{N_{\mathsf{B}}} = 0.8$.

It is easy to check that Eq. 4 is consistent with our analyses in Section IV-B2. In Example 2, $f_{\mathsf{A}} < f_{\mathsf{B}} = f_{\mathsf{C}}$; in Example 3: $f_{\mathsf{B}} = f_{\mathsf{C}} < f_{\mathsf{A}}$.

We define the defection score δ_i of household i as

$$\delta_i = \frac{\kappa(\mathbf{s_{-i}} \cup \omega_i) - \kappa(\mathbf{s})}{e^{o_i}} \tag{5}$$

where $\mathbf{s}$ captures the households' allocations, $\mathbf{s_{-i}} = \mathbf{s} - \{s_i\}$, and $o_i = \frac{|s_i \cap \omega_i|}{v_i} \in [0, 1]$ is the overlap fraction between household i's consumption ω_i and its allocation s_i. $|s_i \cap \omega_i|$ represents the length of the overlapped interval. $o_i = 1$ when household i fully follows the allocation, whereas $o_i = 0$ when household i consumes power fully outside its allocation. For example, if $s_i = (14, 18)$ and $\omega_i = (15, 19)$, then $o_i = \frac{3}{4}$. $\kappa(\mathbf{s_{-i}} \cup \{\omega_i\})$ is the cost to the neighborhood if all households except i follow their allocations. $\kappa(\mathbf{s})$ is the cost to the neighborhood when all households cooperate. In Example 4, $\delta_{\mathsf{A}} = 0$ and $\delta_{\mathsf{B}} > 0$.

Without loss of generality, we normalize the flexibility and defection scores to $[0.5, 1.5]$. Thus the social-cost score Ψ_i for household i is defined as (k is a scaling factor):

$$\Psi_i = k \left(\frac{\delta_i}{\sum_{i \in I} \delta_i} + \frac{1}{2} \right) / \left(\frac{f_i}{\sum_{i \in I} f_i} + \frac{1}{2} \right) \tag{6}$$

Specifically, $f_i > 0$ and $\delta_i = 0$ when the household reports truthfully and $f_i = 0$ and $\delta_i > 0$ when the household misreports and defects. The payment p_i to household i is (ξ is a scaling factor):

$$p_i = \frac{\Psi_i}{\sum_{i \in I} \Psi_i} \cdot \xi \cdot \kappa(\boldsymbol{\omega}) \quad \xi \geq 1 \tag{7}$$

Hence, Enki charges each household the social cost that it burdens the neighborhood without incurring high computational cost [11]. To compute the payment of each agent in a VCG mechanism, one more optimal allocation needs to be calculated assuming the agent was not participating.

A household i's quasilinear utility is the difference between its true valuation and its payment to the neighborhood:

$$U_i(\omega_i | \chi_i, \theta_i) = V_i(\tau_i, v_i, \theta_i) - p_i \tag{8}$$

Enki's payment mechanism incentivizes each household to behave prosocially by (1) linking each household's payment with the neighborhood's cost: a household pays more if the peak is high, and (2) charging each household as much as it harms the neighborhood.

C. Greedy Allocation

With flexibility scores, Enki solves the allocation problem described in Section IV-A in a greedy way. First, Enki computes each household's *predicted* flexibility score, assuming that all of them report truthfully. (Therefore, a household that defects would obtain a positive predicted flexibility score, even though its actual flexibility is zero.) Then Enki proceeds with households in order of increasing predicted flexibility, breaking ties randomly. In Example 3, this mechanism reduces the possibilities: either B or C (picked randomly), then the other, then A. For example, Enki may first allocate $s_{\mathsf{B}} = (18, 20)$ to B. Second, it would tackle C, allocating $s_{\mathsf{C}} = (19, 21)$ to reduce the peak (given that $s_{\mathsf{B}} = (18, 20)$). Third, the allocation for A is $s_{\mathsf{A}} = (16, 18)$.

D. Allocation Objective

Enki always respects each household's preference when it determines an allocation: each household's reported valuation $\hat{V}_i(\tau_i, v_i, \theta_i)$ is ensured to be maximized. In addition, considering the objective of the optimization problem discussed in Section IV-A, Enki's allocation objective, which is the second way Enki aligns each household's self-interest with the group interest, is formalized as

$$\arg\max_{s} \left(\sum_{i \in I} \hat{V}_i(\tau_i, v_i, \theta_i) - \kappa(\boldsymbol{\omega}) \right) \tag{9}$$

Note that when we claim Enki's allocation satisfies *all households' requests*, the requests are *reported* preferences. That's why Enki cannot guarantee strictly Pareto efficiency (Section V-C) even if we make such a claim.

V. Theoretical Analysis

We evaluate the Enki approach in three parts beginning with theoretical analysis. We describe four economic properties [31] and discuss how these properties help identify a good DSM mechanism.

A. Ex Ante Budget Balance

Definition 1 (Ex ante budget balance): A mechanism is ex ante budget balanced if the equilibrium net transfer to the mechanism is not less than zero [11].

Theorem 1: The neighborhood in Enki is ex ante budget balanced. (That is, its revenues are not less than its expenditures.)

Each household i should pay the neighborhood p_i. The utility of the neighborhood U_c is: $U_c = \sum_{i \in I} p_i - \kappa(\boldsymbol{\omega})$. According to Eq. 7, as long as $\xi \geq 1$, Enki can ensure the neighborhood is ex ante budget balanced:

$$U_c = \sum_{i \in I} p_i - \kappa(\boldsymbol{\omega}) = (\xi - 1)\kappa(\boldsymbol{\omega}) \geq 0.$$

B. Bayesian Incentive Compatibility

Like VCG, Enki is a mechanism in a Bayesian game setting, where information about characteristics of other players (i.e., payoffs) is incomplete. However, we relax the definition of incentive compatibility from dominant strategies to Bayes-Nash equilibrium strategies. A Bayes-Nash equilibrium is a Nash equilibrium in a Bayesian game [14].

Definition 2 (Bayesian incentive-compatibility): A mechanism is Bayesian incentive-compatible if, for any individual in Bayes-Nash equilibrium, truth-telling is the best strategy, or at least not worse than another strategy [14].

Enki is Bayesian incentive-compatible if each household obtains the most utility by truthfully reporting its preferences, given that all the others report truthfully. We further adopt the idea of weak Bayesian incentive-compatibility from Rose et al. [24], which relaxes incentive compatibility so that it needs to hold only in expectation for each household. We can show that Enki achieves weak Bayesian incentive compatibility.

Theorem 2: Enki is weakly Bayesian incentive-compatible.

We use an example with two scenarios to motivate Bayesian incentive-compatibility. In the first scenario, household A with a true preference of $\chi_{\mathsf{A}} = (18, 20, 2)$ misreports its preference as $\hat{\chi_{\mathsf{A}}} = (14, 20, 2)$, and obtains an allocation $s_{\mathsf{A}} = (14, 16)$. Instead, A defects and consumes power during $\omega_{\mathsf{A}} = (18, 20)$. In this case, for the same total consumption, A should have provided $\hat{\chi_{\mathsf{A}}} = (18, 20, 2)$ rather than $\hat{\chi_{\mathsf{A}}} = (14, 20, 2)$. In the second scenario, A provides us with its true interval. Thus, Bayesian incentive-compatibility holds only if A's utility in the second scenario is higher. For simplicity, we omit some parameters.

Following Eq. 8, A's utility in the first scenario is:

$$U'_{\mathsf{A}} = V_a(\tau'_a, v_a, \theta_a) - \xi \cdot \kappa(\boldsymbol{\omega}') \frac{\Psi'_a}{\sum_{i \in I} \Psi'_i}$$

In the second scenario, A earns a utility as follows:

$$U''_{\mathsf{A}} = V_a(\tau''_a, v_a, \theta_a) - \xi \cdot \kappa(\boldsymbol{\omega}'') \frac{\Psi''_a}{\sum_{i \in I} \Psi''_i}$$

Specifically, $\Psi'_a = 2k \cdot \Delta_a = 2k \left(\frac{\delta_a}{\sum_{i \in I} \delta_i} + \frac{1}{2} \right)$ and $\Psi''_a = \frac{k}{2} \frac{1}{F_a} = \frac{k}{2} \left(\frac{1}{\frac{f_a}{\sum_{i \in I} f_i} + \frac{1}{2}} \right)$, where Δ_a and F_a are in $[0.5, 1.5]$. Hence, $F_a \cdot \Delta_a \geq \frac{1}{4}$. Thus, $2k \cdot \Delta_a \geq k/(2F_a)$. Hence, we obtain $\Psi'_a \geq \Psi''_a$.

In addition, $\tau'_a \leq v_a$ because A misreports its preference in the first scenario, whereas $\tau''_a = v_a$ in the second scenario, $V_a(\tau'_a, v_a, \theta_a) \leq V_a(\tau''_a, v_a, \theta_a)$. Since consumptions are equal in both scenarios, i.e., $\boldsymbol{\omega}' = \boldsymbol{\omega}''$, we see that $\kappa(\boldsymbol{\omega}') = \kappa(\boldsymbol{\omega}'')$.

Next, we check the relationship between $\Psi'_a / \sum_{i \in I} \Psi'_i$ and $\Psi''_a / \sum_{i \in I} \Psi''_i$, when all households, except A, have submitted their true preferences. First, their defection scores are zero $\forall i \neq a, \delta_i = 0$. Second, A's flexibility score f_a is zero in the first scenario, whereas it becomes positive in the second scenario. Because the sum of normalized flexibility scores ($\sum_{i \in I} F_i$) is fixed, the average normalized flexibility score of others decreases in the second scenario.

Based on the law of large numbers theorem [32], which states that the sample mean converges to the distribution mean as the sample size increases, the expected normalized flexibility score of others decreases when the number of households becomes large ($\forall i \neq a, \bar{F}'_i \geq \bar{F}''_i$). Accordingly, their expected social-cost scores (Eq. 6) in the second scenario are increasing ($\forall i \neq a, \bar{\Psi}'_i \leq \bar{\Psi}''_i$). Therefore, $\Psi'_a / \sum_{i \in I} \Psi'_i$ is not less than $\Psi''_a / \sum_{i \in I} \Psi''_i$.

Combined with the previous analysis, $U'_{\mathsf{A}} \leq U''_{\mathsf{A}}$ when the number of households is large. However, Enki may not be strictly incentive-compatible when the number of households is small. Hence, Enki is weakly Bayesian incentive-compatible.

Enki could be made Bayesian incentive-compatible by setting the payment of each household i as $p_i = \Psi_i \cdot \kappa(\boldsymbol{\omega})$.

Instead, in order to ensure budget-balance, we set the payment function as shown in Eq. 7, thereby limiting Enki to weakly Bayesian incentive-compatibility. This is as expected: a VCG mechanism cannot respect budget-balance if it is efficient and strategy proof [11].

C. Pareto Efficiency

Definition 3 (Pareto efficiency): A mechanism is Pareto efficient if its choice in equilibrium ensures that the sum of the *true* valuations from all agents is maximized [11].

We can formalize the definition of Pareto efficiency in Enki as $\forall V \forall \mathbf{s}', \sum_{i \in I} V_i(\mathbf{s_i}) \geq \sum_{i \in I} V_i(\mathbf{s_i}')$, where $V_i(\mathbf{s_i})$ is equal to $V_i(\tau_i, v_i, \theta_i)$.

Theorem 3: Enki is weakly Pareto efficient.

Enki's allocation objective maximizes $\sum_{i \in I} \hat{V}_i(\mathbf{s_i}) - \kappa(\boldsymbol{\omega})$ (Eq. 9). And, we have proved that Enki is weakly incentive-compatible in Section V-B, that is, each household would reveal its true preference in equilibrium. Therefore, the following is maximized in equilibrium: $\sum_{i \in I} V_i(\mathbf{s_i}) - \kappa(\boldsymbol{\omega})$. Further, $\kappa(\boldsymbol{\omega})$ is nearly minimized under Enki's greedy allocation. Thus, $\sum_{i \in I} V_i(\mathbf{s_i})$ is nearly maximized in equilibrium, and Enki is accordingly weakly Pareto efficient.

D. Individual Rationality

Definition 4 (Individual rationality): A mechanism is individually rational if each agent would receive a nonnegative utility from participating in the mechanism [33].

Theorem 4: Enki is not individual rational.

As for a VCG mechanism, individual rationality cannot be ensured in Enki. In the quasilinear setting, a household's valuation is independent of its payment. Once a given household has reported its true preference, its valuation stays the same. However, the household's payment would still increase since its payment depends on the peak of power usage. Thus, a household sometimes would receive a negative utility. This conclusion accords with the Myerson-Satterthwaite theorem [33], which states that no mechanism exists that is ex ante budget-balanced, efficient, Bayesian incentive-compatible, and individually rational.

However, Enki still provides incentives for households, especially those flexible ones, to participate.

Theorem 5: The expected utility of all households is higher with Enki ($\mathbb{E}(U_i)$).

In contrast with traditional application of mechanism design, a household's utility should not be zero when it does not participate in Enki, because every household should consume power. When not participating, we define every household as a *price taking* user: it doesn't know the effect of its actions on the price, and therefore on its payment; it considers that the price rate is set by the power company, even though it is not true. (In fact, a household's consumption will affect the aggregate consumption as well as its own payment; we assume that households disregard such a connection). The utility of each household is the same as we defined in Eq. 8, whereas the payment mechanism is a little different. We use the *proportional allocation mechanism* [34]. That is, each household pays a price that is proportional to its power usage instead of its social cost: $p_i^z = \frac{b_i}{\sum_{i \in I} b_i} \cdot \xi \cdot \kappa(\boldsymbol{\omega^z})$, where b_i is the amount of power consumption, and $\xi \geq 1$. In addition, the valuation of each household stays the same no matter whether it participates in Enki, because each household's preference is respected in Enki's allocation.

The cost of the neighborhood without Enki ($\kappa(\boldsymbol{\omega^z})$) should be higher, because the greedy allocation in Enki reduces the peak. Let n represent the number of households. We omit some parameters to simplify the notation.

$$\mathbb{E}(U_i) = \left(\sum_{i \in I} V_i - \xi \cdot \kappa(\boldsymbol{\omega}) \right) / n \geq \left(\sum_{i \in I} V_i - \xi \cdot \kappa(\boldsymbol{\omega^z}) \right) / n$$
$$= \sum_{i \in I} \left(V_i - \frac{b_i}{\sum_{i \in I} b_i} \cdot \xi \cdot \kappa(\boldsymbol{\omega^z}) \right) / n = \mathbb{E}(U_i^z)$$

Theorem 6: The expected utility of a flexible household is higher with Enki ($\mathbb{E}(U_f)$).

Without losing generality, let's assume that all the households consume the same amount of power while household F is most flexible, leading to a least social cost of household F.

$$\begin{aligned} \mathbb{E}(U_f) &= V_f - \frac{\Psi_f}{\sum_{i \in I} \Psi_i} \cdot \xi \cdot \kappa(\boldsymbol{\omega}) \geq V_f - \frac{1}{n} \cdot \xi \cdot \kappa(\boldsymbol{\omega}) \\ &\geq V_f - \frac{1}{n} \cdot \xi \cdot \kappa(\boldsymbol{\omega^z}) = V_f - \frac{b_f}{\sum_{i \in I} b_i} \cdot \xi \cdot \kappa(\boldsymbol{\omega^z}) \\ &= \mathbb{E}(U_f^z) \end{aligned}$$

Overall, even if Enki is not individually rational according to the traditional definition, households behave rationally by participating in Enki, especially those with higher flexibility.

VI. Simulation Study

Because the performance of Enki's greedy allocation hasn't been studied yet, and Enki is weakly incentive-compatible, we carry out a simulation study to address three questions. First, how effective is Enki in reducing peak consumption and thus reducing the cost to the neighborhood? Second, how much faster is Enki's greedy scheduler than the optimal scheduler? Third, how effective is Enki in incentivizing each household to reveal its true preference?

In our simulation, each household has a *usage profile*: a *narrow interval*, a *wide interval*, and a *duration*. A household *most prefers* to consume power during its narrow interval, whereas it *can* consume power during its wide interval. For example, A household's narrow interval to watch TV is 7PM to 8PM, and it can watch TV for one hour during the wide interval between 6PM and 11PM without affecting its comfort too much. But the household would not like to watch TV before 6PM or after 11PM.

A Poisson distribution with mean 16 generates the beginning times of the narrow and wide intervals. A uniform distribution [1, 4] generates the duration. The ending time of the narrow interval is *beginning time* + *duration*. The wide interval's ending time comes from a uniform distribution [*ending time of its narrow interval*+2, 24]. The power consumption is 2 kWh.

The valuation factor ρ comes from a uniform distribution [1, 10]. Scaling factors are set to $\sigma = 0.3$, $k = 1$, $\xi = 1.2$.

A. Social Welfare

The MIQP solver in IBM ILOG CPLEX V12.4 implements Optimal (Eq. 2). Our metrics are the *peak to average ratio* (PAR) and the cost to the neighborhood. We consider a neighborhood with a population ranging from 10 to 50. For each population, the simulation repeats 10 rounds to simulate 10 days. Further, every household reports its wide interval as its true preference, and its preference is generated at the beginning of every day. Results are averaged and the plotted error bars represent the 95% confidence intervals. Figures 4 and 5 show that differences between the PARs and between the neighborhood's costs of the two allocations are not large.

Nevertheless, the time the two allocations take for scheduling is significantly different (Figure 6). When the number of households is over 40, Optimal on average takes around 600 times longer.

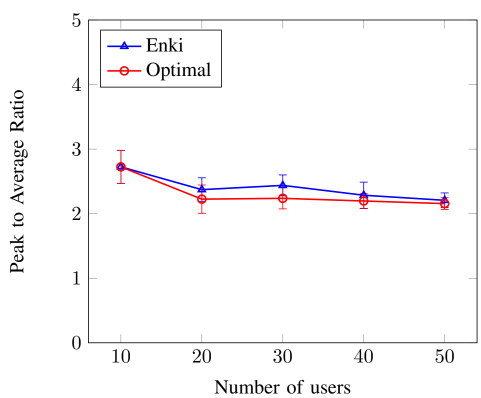

Fig. 4: Peak to average ratio (PAR) for Enki and Optimal.

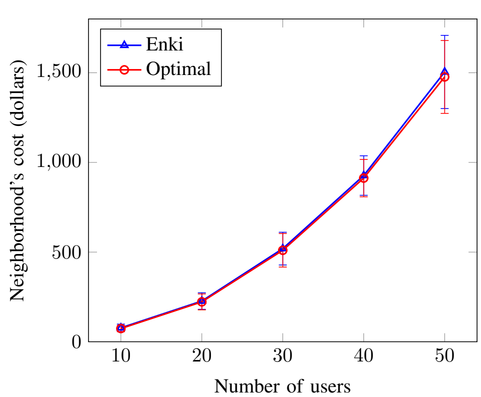

Fig. 5: Cost to the neighborhood for Enki and Optimal.

B. Incentive Compatibility

We consider a neighborhood with $N = 50$ households. For the first household, we assume its narrow interval is $(18, 20)$ while its wide interval is $(16, 24)$; its true preference is its narrow interval, and its valuation factor is five. For the other

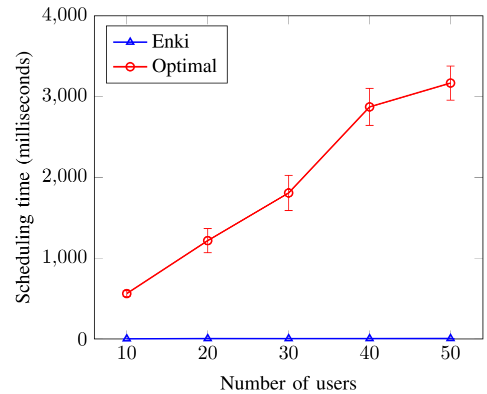

Fig. 6: Time taken to compute Enki's allocation and Optimal.

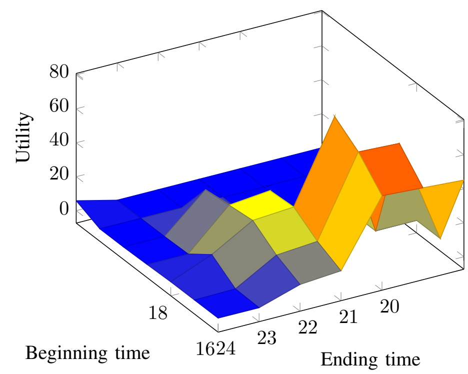

Fig. 7: Utility of the first household when all other households report their preferences truthfully.

households, we assign a narrow interval to each of them as their true preferences; we generate their usage profiles at the beginning of the first day and keep them unchanged.

We explore the best response of the first household when the other households declare their preference truthfully. This process repeats 10 times, and Figure 7 shows the average utility of the first household for each possible preference that it could report. The first household's best response happens when it declares its true preference $(18, 20)$; Enki is weakly incentive compatible.

VII. User Study

The objective of Enki is, in essence, to elicit the true valuation from each household as a way of reducing peak consumption. However, Enki's mechanism is based on two assumptions. One, "rationality" assumes that each individual behave in a way to maximize his or her utility. Two, "consistency" ignores the diversity of human subjects. To verify whether Enki's objective is maintained in the real world, we conduct a user study via an online game. Evaluating Enki in its entirety is beyond the score of this paper; we instead design the user study to investigate two research questions:

- RQ_1, how effective is Enki in preventing subjects from defecting?

- RQ$_2$, how effective is Enki in incentivizing subjects to reveal a wider true interval?

A. Subjects

We recruited 20 computer science students (four female; three undergraduates; four with prior gambling experience). None had prior experience with economic experiments. The idea behind this population of subjects was that they are surrogates for ordinary consumers who have limited knowledge of power systems. However, we recognize that they are not like ordinary consumers with regard to other attributes, such as being college educated and technically savvy. This remains a threat to the validity of our empirical evaluation results.

B. Game Information

The study had three stages. In the instruction stage, we explained the game (as below) and its objective to maximize points earned. Each subject represents a household and interacts with Enki (embedded in the online game) to simulate a neighborhood. The game playing stage includes 16 rounds, each simulating one day. We provided each subject a true preference (a true interval and a duration), and followed these steps:

- *Preference*. Each subject reports a preferred interval. In the real application, the report will be done through ECC automatically. In the study, we assume that each subject has perfect information of his power demand for the next day.
- *Allocation*. Enki assigns an allocation to each subject.
- *Household consumption*. We automate this step: selecting real consumption to be within the subject's true interval and close to his allocation.
- *Household payment*. Each subject pays the center based on his submission and his true preference.
- *Calculation*. We calculate each subject's utility following Eq. 8 and transform each subject's utility into a score between zero and 100.
- *Information*. We reveal each subject's load statistics (his consumption versus other subjects' consumption) and score history to himself.

In a post-study questionnaire, we asked our subjects about their background information, understanding, and risk attitudes [11] of playing the game.

Importantly, we tempt subjects to perform unsocial behaviors (misreporting): a subject may potentially earn a higher score by broadening or shifting his or her submitted interval. But subjects are informed that they may lose points by defection. To reduce complexity, we provide subjects a calculator to help them estimate their payoffs from different intervals before they submit an interval [35].

We paid subjects between \$10 and \$20, depending on their performance: a fixed payment of \$10 for participating plus one cent for each point they earned.

C. Study Setting

The experiment has two treatments: each treatment has four sessions. Each session takes approximately 40 minutes. Treatment 1 has a group of subjects while Treatment 2 has only one subject. To control the experiment, we add six artificial agents in Treatment 1 and four artificial agents in Treatment 2. The true preference of each subject changes every four rounds so that they can learn the game and adjust their strategies. For each artificial agent, its true preference updates every round. Half of the agents defect in Rounds 1 to 8 whereas all agents cooperate in Rounds 9 to 16.

D. Results

RQ$_1$: We define the "stages" Overall, Initial, Defect, and Cooperate respectively as Rounds 1 to 16, 1 to 4, 1 to 8, and 9 to 16. In each stage, a subject's *defection rate* is the number of rounds the subject defects divided by the stage's number of rounds.

Table II shows that (1) the average defection rate of all subjects in Overall is low; (2) subjects tend to defect more often in Initial, while learning the game and tend to less often in Cooperate, where all artificial agents cooperate. This corroborates that Enki is weakly incentive-compatible.

TABLE II: Average Defection Rate of 20 Subjects

Overall	Initial	Defect	Cooperate
0.2049	0.3625	0.2938	0.125

We conduct a Mann-Whitney U test [36] to test the above observations. Each stage has two samples, each sample being a vector of size 20. Sample 1 contains the number of rounds that a subject defects; Sample 2 assumes that each subject defects randomly, thus the value of each element is half of the number of rounds in that particular stage. For example, Overall covers 16 rounds, so the value of each element in Sample 2 is eight.

In Table III, the p-value in Overall is less than 0.0001: the means between the two samples is significantly different, and thus **Enki is effective in preventing subjects from defecting**. Further, the difference is not significant in Initial but is so in Defect and Cooperate.

To compare the two treatments' effect on performance, we make two groups, one being the 16 subjects in Treatment 1 and the other being the four subjects in Treatment 2. We calculate the average defection rate within each group in different stages. Table IV shows that subjects in Treatment 2 tend to defect less often, especially during Cooperate. This corroborates the claim that Enki is weakly incentive-compatible: each subject in Treatment 2 makes a decision when all the other subjects (artificial agents) cooperate during Cooperate.

RQ$_2$: Results regarding RQ$_1$ show that subjects learn the game during Initial and gain familiarity with it during Cooperate; if a subject has understood the game, he will choose exactly his true interval more often in Cooperate than in Initial. *True interval selecting ratio* is the ratio of the number of rounds a subject chooses his exact true interval in one

TABLE III: Mann-Whitney U Test of Whether Enki is Effective in Preventing Defection

	Overall	Initial	Defect	Cooperate
Each element in Sample 1	# rounds that a subject defects			
Each element in Sample 2	8	2	4	4
p-value	< 0.0001	0.0532	0.0078	< 0.0001

TABLE IV: Average Defection Rate in Two Treatments

	Overall	Initial	Defect	Cooperate
T1	0.23	0.34	0.31	0.15
T2	0.14	0.44	0.25	0.03

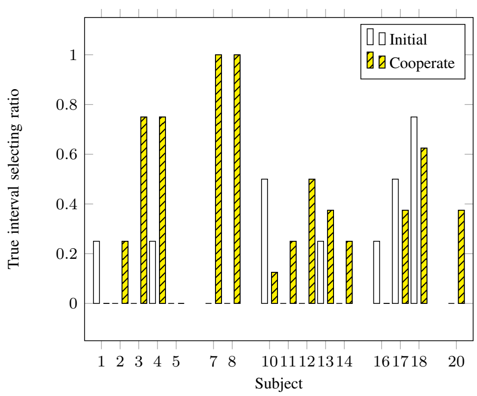

Fig. 8: Mann-Whitney U test of whether subjects select their true interval more often in Cooperate than in Initial.

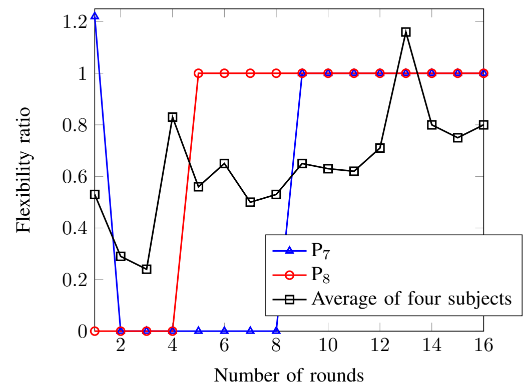

Fig. 9: Flexibility for typical subjects and two subjects who understand well.

stage to the number of rounds in that stage. The average true interval selecting ratio of 20 subjects in Initial is 23.75% and it increases to 37.5% in Cooperate.

Four subjects reported in the post-study questionnaire that they had not understood the game at all: they randomly submitted an interval in each round. We removed the data of these four subjects and conducted a Mann-Whitney U test for the remaining 16 subjects. Sample 1 contains the true interval selecting ratio of each subject in Initial and Sample 2 contains the ratio in Cooperate, as shown in Figure 8. The p-value is 0.0143 supports the hypothesis that subjects tend to submit their true interval more often in Cooperate than in Initial.

In addition, we define a subject's *flexibility ratio* as the ratio of the length of the submitted interval within the true interval and the length of the true interval. This ratio ranges from zero (subject defects) to one (subject chooses his exact true interval). Interestingly, by observing changes in the flexibility ratio for each subject, we find that two subjects (P_7 and P_8) who reported that they understood the game well share a similar changing pattern (Figure 9). Both of them defect often in Initial and then stick to their exact true interval in Cooperate. Figure 9 also shows that the average flexibility ratio of four subjects, who reported an intermediate understanding of the game, increases. That is, we can conclude that **subjects are incentivized to reveal a wider true interval in Enki**. The result additionally indicates the importance of developing intuitive user interfaces, so that consumers would properly understand the Enki mechanism.

VIII. Conclusions

We propose Enki as a novel tractable mechanism that reduces peak total consumption and encourages prosocial behavior through its social objective and payment mechanism. We show via theoretical analysis that Enki is ex-ante budget-balanced, weakly Pareto-efficient, and weakly Bayesian incentive-compatible. A simulation study verifies that Enki approximates the optimal allocation strategy while substantially reducing computation time in producing an allocation that is acceptable to all participants. A major challenge with any new economic mechanism is that if its target users fail to understand it and confidently participate in it, the mechanism will fail no matter how theoretically powerful it is. A human-subject study yields positive results by showing that Enki motivates subjects to avoid defection and is effective in incentivizing them to reveal flexible preferences.

In future work, we will investigate a decentralized mechanism and consider direct cooperation among households forming small coalitions to reduce their joint peak demand further. We will model more complex user behaviors than at present and a variety of appliances [37]. In particular, we are interested in approaches that not only reduce peak demand but reduce aggregate demand (i.e., save power not just shift load).

Acknowledgments

Thanks to Mesut Baran, Mona Singh, and Neeraj Saxena for useful discussions and to the anonymous reviewers for helpful comments. This research was partially supported by the NCSU Laboratory for Analytic Sciences and by the National Science Foundation under grant 0910868.

REFERENCES

[1] S. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings, "Putting the "smarts" into the Smart Grid: A grand challenge for artificial intelligence," *Communications of the ACM*, vol. 55, no. 4, pp. 86–97, 2012.

[2] SPP, "Impact evaluation of the California statewide pricing pilot," 2005.

[3] D. Challet, M. Marsili, and Y.-C. Zhang, *Minority Games: Interacting Agents in Financial Markets*, ser. Oxford Finance Series. Oxford: Oxford University Press, 2005.

[4] Y. Guo, M. Pan, Y. Fang, and P. P. Khargonekar, "Decentralized coordination of energy utilization for residential households in the smart grid," *IEEE Transactions on Smart Grid*, vol. 4, no. 3, pp. 1341–1350, Sept 2013.

[5] P. Chavali, P. Yang, and A. Nehorai, "A distributed algorithm of appliance scheduling for home energy management system," *IEEE Transactions on Smart Grid*, vol. 5, no. 1, pp. 282–290, 2014.

[6] A.-H. Mohsenian-Rad, V. W. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, "Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future Smart Grid," *IEEE Transactions on Smart Grid*, vol. 1, no. 3, pp. 320–331, Dec. 2010.

[7] H. Chen, Y. Li, R. H. Louie, and B. Vucetic, "Autonomous demand side management based on energy consumption scheduling and instantaneous load billing: An aggregative game approach," *IEEE Transactions on Smart Grid*, vol. 5, no. 4, pp. 1744–1754, 2014.

[8] W. Tushar, W. Saad, H. V. Poor, and D. B. Smith, "Economics of electric vehicle charging: A game theoretic approach," *IEEE Transactions on Smart Grid*, vol. PP, no. 99, pp. 1–12, 2012.

[9] P. Samadi, H. Mohsenian-Rad, R. Schober, and V. W. S. Wong, "Advanced demand side management for the future Smart Grid using mechanism design," *IEEE Transactions on Smart Grid*, vol. 3, no. 3, pp. 1170–1180, Sep. 2012.

[10] C. A. Goldman, N. C. Hopper, O. Sezgen, M. M. Moezzi, R. Bharvirkar, B. Neenan, R. Boisvert, P. Cappers, and D. Pratt, "Customer response to day-ahead wholesale market electricity prices: Case study of RTP program experience in New York," U.S. Department of Energy, LBNL, Tech. Rep., 2004.

[11] Y. Shoham and K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, Dec. 2008.

[12] N. Nisan and A. Ronen, "Algorithmic mechanism design (extended abstract)," in *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, ser. STOC '99. ACM, 1999, pp. 129–140.

[13] S. K. Barker, A. K. Mishra, D. E. Irwin, P. J. Shenoy, and J. R. Albrecht, "Smartcap: Flattening peak electricity demand in smart homes," in *PerCom*. IEEE Computer Society, 2012, pp. 67–75.

[14] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. New York, NY, USA: Cambridge University Press, 2007.

[15] Y. Guo, R. Li, G. Poulton, and A. Zeman, "A simulator for self-adaptive energy demand management," in *Proceedings of the International IEEE Conference on Self-Adaptive and Self-Organizing Systems (SASO)*, Oct. 2008, pp. 64–73.

[16] S. J. Darby, "Load management at home: advantages and drawbacks of some 'active demand side' options," *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 2012.

[17] A. Faruqui and S. Sergici, "Household response to dynamic pricing of electricity: A survey of 15 experiments," *Journal of Regulatory Economics*, vol. 38, pp. 193–225, 2010.

[18] G. Strbac, "Demand side management: Benefits and challenges," *Energy Policy*, vol. 36, no. 12, pp. 4419–4426, 2008.

[19] N. Gatsis and G. B. Giannakis, "Cooperative multi-residence demand response scheduling," in *Proceedings of the 45th Annual Conference on Information Sciences and Systems (CISS)*, Mar. 2011, pp. 1–6.

[20] T. T. Kim and H. V. Poor, "Scheduling power consumption with price uncertainty," *IEEE Transactions on Smart Grid*, vol. 2, no. 3, pp. 519–527, Sep. 2011.

[21] A.-H. Mohsenian-Rad, V. W. Wong, J. Jatskevich, and R. Schober, "Optimal and autonomous incentive-based energy consumption scheduling algorithm for Smart Grid," in *Innovative Smart Grid Technologies (ISGT)*, Jan. 2010, pp. 1–6.

[22] S. Salinas, M. Li, and P. Li, "Multi-objective optimal energy consumption scheduling in smart grids," *IEEE Transactions on Smart Grid*, vol. PP, no. 99, pp. 1–8, 2012.

[23] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. Jennings, "Agent-based control for decentralised demand side management in the Smart Grid," in *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2011, pp. 5–12.

[24] H. Rose, A. Rogers, and E. H. Gerding, "A scoring rule-based mechanism for aggregate demand prediction in the Smart Grid," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Jun. 2012.

[25] S. Stein, E. Gerding, V. Robu, and N. Jennings, "A model-based online mechanism with pre-commitment and its application to electric vehicle charging," in *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2012, pp. 669–676.

[26] Z. Zhou, F. Liu, Z. Li, and H. Jin, "When smart grid meets geo-distributed cloud: An auction approach to datacenter demand response." in *INFOCOM*. IEEE, 2015, pp. 2650–2658.

[27] Q. Sun, S. Ren, C. Wu, and Z. Li, "An online incentive mechanism for emergency demand response in geo-distributed colocation data centers," in *Proceedings of the Seventh International Conference on Future Energy Systems*. ACM, 2016, pp. 3:1–3:13.

[28] B. Schneier, *Liars and Outliers: Enabling the Trust that Society Needs to Thrive*. Indianapolis: Wiley, 2012.

[29] G. Hardin, "The tragedy of the commons," *Science*, vol. 162, no. 3859, pp. 1243–1248, Dec. 1968.

[30] W. Tushar, W. Saad, H. V. Poor, and D. B. Smith, "Economics of electric vehicle charging: A game theoretic approach," *CoRR*, vol. abs/1208.0631, 2012.

[31] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*. Oxford University Press, 1995.

[32] W. Feller, *An Introduction to Probability Theory and Its Applications*. Wiley, 1968, vol. 1.

[33] R. Myerson and M. Satterthwaite, "Efficient mechanisms for bilateral trading," *Journal of Economic Theory*, vol. 29 (1983), pp. 265–281, 1981.

[34] F. P. Kelly, "Charging and rate control for elastic traffic," *European Transactions on Telecommunications*, vol. 8, pp. 33–37, 1997.

[35] Y. Masatlioglu, S. Taylor, and N. Uler, "Behavioral mechanism design: evidence from the modified first-price auctions," *Review of Economic Design*, vol. 16, no. 2-3, pp. 159–173, 2012.

[36] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," *The Annals of Mathematical Statistics*, vol. 18, no. 1, pp. 50–60, 1947.

[37] B. Aksanli, A. S. Akyurek, and T. S. Rosing, "User behavior modeling for estimating residential energy consumption," in *Proceedings of EAI International Conference on Smart Grids for Smart Cities*. Springer, 2015.