

Heuristic Optimization for an Aggregator-Based Resource Allocation in the Smart Grid

Timothy M. Hansen, *Graduate Student Member, IEEE*, Robin Roche, *Member, IEEE*,
Siddharth Suryanarayanan, *Senior Member, IEEE*, Anthony A. Maciejewski, *Fellow, IEEE*,
and Howard Jay Siegel, *Fellow, IEEE*

Abstract—We utilize a for-profit aggregator-based residential demand response (DR) approach to the smart grid resource allocation problem. The aggregator entity, using a given set of schedulable residential customer assets (e.g., smart appliances), must set a schedule to optimize for a given objective. Here, we consider optimizing for the profit of the aggregator. To encourage customer participation in the residential DR program, a new pricing structure named customer incentive pricing (CIP) is proposed. The aggregator profit is optimized using a proposed heuristic framework, implemented in the form of a genetic algorithm, that must determine a schedule of customer assets and the CIP. To validate our heuristic framework, we simulate the optimization of a large-scale system consisting of 5555 residential customer households and 56 642 schedulable assets using real-pricing data over a period of 24-h. We show that by optimizing purely for economic reasons, the aggregator can enact a beneficial change on the load profile of the overall power system.

Index Terms—Aggregator, appliance scheduling, customer incentive pricing (CIP), cyber-physical systems (CPSs), heuristic optimization, smart grid.

I. INTRODUCTION

ACCORDING to the U.S. Department of Energy, since 1982, the growth in peak electricity usage has exceeded the growth in transmission capacity by almost 25% each year [1]. Furthermore, electricity sales in the residential sector in the U.S. are expected to grow 24% from the 2011 reference case to 2040 [2]. Given these trends, peak energy demands are expected to exceed the available transmission capability. This can be dealt with by increasing transmission capability, creating distributed generation (DG), or curtailing load. As shown in [1], it is unlikely that additional spending will be allocated for increasing transmission capability, leading to research in the areas of DG and, in the case of this paper, curtailing load during peak hours.

Manuscript received November 26, 2013; revised May 19, 2014, August 12, 2014, and December 8, 2014; accepted January 23, 2015. Date of publication February 20, 2015; date of current version June 18, 2015. This work was supported in part by the National Science Foundation under Grant CNS-0905399 and Grant CCF-1302693, in part by Colorado State University G. T. Abell Endowment, and in part by the Université de Technologie de Belfort-Montbéliard, France. Paper no. TSG-00879-2013.

T. M. Hansen, S. Suryanarayanan, A. A. Maciejewski, and H. J. Siegel are with Colorado State University, Fort Collins, CO 80523 USA (e-mail: timothy.hansen@colostate.edu).

R. Roche is with the Research Institute on Transportation, Energy, and Society, Université de Technologie de Belfort-Montbéliard, Belfort 90010, France.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSG.2015.2399359

In addition to the physical considerations, there is also an economical motivation. By curtailing load during peaks, electricity costs could be drastically reduced by eliminating the need for peaking power plants. From [3], “a 5% reduction in peak demand during the California energy crisis of 2000–2001 would have reduced the highest wholesale prices by 50%.” We attempt to reduce peak demands by intelligently coordinating the scheduling of customer appliances away from the peak time, alleviating the peak demand, and offering a benefit to all parties.

Given both the physical and economical motivations, an aggregator-based residential demand response (DR) program is presented. The aggregator is a proposed for-profit entity in a deregulated market structure that interfaces a DR market (DRX) and a set of customers. The aggregator will possess information about the schedulable assets of the participating customers. In many, if not all energy markets, there is a minimum power rating required to bid into the market (e.g., 0.1 MW in the PJM market [4]). By aggregating the customer assets, the aggregator is able to enact a noticeable change on the overall system by scheduling the assets of many customers and bidding the aggregation of their assets, where a single customer would not be able to do so.

This provides the customer the opportunity to participate in the deregulated electricity market, through the aggregator. As presented in [5], incentives can influence customer behavioral changes. To encourage customers to participate with the aggregator on a daily basis, we propose a new day-ahead price for electricity offered by the aggregator, in the form of customer-incentive pricing (CIP), to offset the customers’ inconvenience of the aggregator controlling their assets. Additionally, if the inconvenience of rescheduling the load is not worth the reduced price, the customer may refuse the aggregator and instead pay the utility company for electricity.

The aggregator-based residential DR program, denoted smart grid resource allocation (SGRA), is formally stated as given a set of customers and information about their respective assets, subject to customer constraints (i.e., availability of customer assets and customer incentive requirements), how can the aggregator find the CIP and schedule of assets to maximize aggregator profit? We demonstrate that by optimizing solely for the profit of the aggregator, we can enact a change on the peak load of the system because this is where most of the profit can be made due to the high cost of peak generators.

To solve the SGRA problem, we borrow concepts from resource allocation in computing where tasks must be allocated to machines to optimize a performance metric, such as completing all tasks as quickly as possible. It has been shown, in general, that such problems are NP-complete [6]–[8] and, as such, use heuristic optimization to find near-optimal solutions. Similarly in this paper, we use heuristic optimization techniques to find near-optimal solutions to the SGRA problem in a time frame that is reasonable with the large number of assets considered for use as a day-ahead scheduler.

Related prior work on demand side management in smart grid has occurred in the areas of optimization and aggregation of end-user resources. The optimization of scheduling end-user resources has been approached as linear programming [9], [10], dynamic programming [11], and mixed integer programming [12]. Heuristic-based methods also have been used in the form of particle swarm optimization [13], evolutionary algorithms [14], and multiagent systems [15]. However, increasing DR technologies and allowing retail customers direct access to wholesale market prices may increase the price-elasticity of demand, leading to increased volatility in power systems [16]. Aggregators are an intermediary entity that offer centralized coordination of many entities [17]–[20]. This paper differs in that we schedule many more distributed residential customer assets and introduce a new time-variant customer pricing mechanism in the form of CIP, existing in conjunction with the utility price, to encourage customer participation. Active distribution networks with full integration of demand and distributed energy resources is an aggregator framework that benefits in a market from the active participation of residential and commercial consumers [21], [22]. Our framework differs from [21] and [22] in our market approach, specifically CIP, and the aggregator-customer-utility relationship.

According to the California Energy Commission (CEC), residential loads are not easily controlled and need to be composed of a large portfolio of assets to provide a strategic DR product [23]. The CEC identified strategies to fulfill its DR requirements including: direct DR participation with the independent system operator (ISO), new market and auction mechanisms (e.g., our proposed DRX), improving customer willingness to participate, and the introduction of time-variant pricing. This paper directly addresses each of these strategies, offering direct DR participation through the customer-aggregator-DRX relationship (see Fig. 1) and encouraging customer participation with the time-variant CIP mechanism.

In this paper, we theorize that by using an aggregator placed between the customer and bulk power market the volatility in the power system can be reduced. In [24], we hypothesized the use of a heuristic approach to the SGRA. In this paper, we design and implement the heuristic framework using a simulation test bed of 5555 customers to simulate the scheduling of over 56 000 consumer devices centrally controlled by the aggregator. In this paper, we make the following unique contributions.

- 1) A new customer pricing structure is proposed in the form of CIP to encourage customer participation in residential DR.

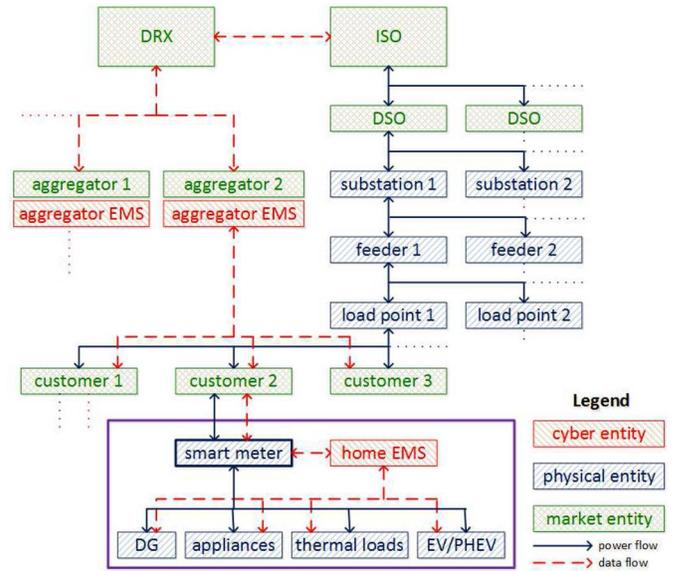


Fig. 1. Architecture and communication for the cyber-physical system (CPS) of the proposed aggregator-based residential DR program.

- 2) A heuristic optimization framework is designed to implement and solve the SGRA problem.
- 3) An analysis of our heuristic framework using actual electricity pricing data and a large-scale simulation test system consisting of 5555 customers and 56 642 schedulable assets is conducted.
- 4) An aggregator-based approach for a residential DR program for use in scheduling customer assets in a large-scale manner.

In our simulation study, by optimizing for profit, the aggregator was able to reduce the peak load of the 5555 participating customers by 12.5%. We demonstrate that this change benefits the customer of the aggregator (in the form of reduced cost of electricity for schedulable loads), the aggregator (in the form of a profit), and also those customers not participating with the aggregator (because the overall system peak is lowered as a common good).

The rest of this paper is organized as follows. Section II describes the system model and the enabling technologies. In Section III, a heuristic framework and genetic algorithm (GA) implementation are presented. The setup for the simulation study is discussed in Section IV. Section V examines the simulation results. Section VI concludes this paper.

II. SYSTEM MODEL

A. CPS

The proposed CPS for the aggregator-based residential DR program is shown in Fig. 1. On the right of Fig. 1 is the traditional power system and market structure that flows from the ISO to the distribution system operator for delivering electricity to the residential customer. The left-hand side of Fig. 1 encapsulates our proposed residential DR program. The DRX is an ancillary market in a fully deregulated market structure that provides DR services to the ISO. The aggregator interfaces the DRX and the residential customer, and provides the

positive attributes (e.g., load shifting and distributed storage) of the aggregated customer assets (e.g., DG and electric vehicles) to the ISO. Each participating customer has a home energy management system (HEMS) that controls the assets, connected to a smart meter. The aggregator coordinates the use of the participating customer assets and brings the result (e.g., load reduction) to the DRX for market exchange. The aggregator and customer interactions will be expanded on in the following subsections.

For realizing the market interactions in Fig. 1, several enabling technologies are first expected to penetrate the electric power system, and are assumed to exist in this paper. As previously mentioned, the retail electricity market must be fully deregulated, allowing for the customer to choose between suppliers. The control and communication infrastructure, including the requisite cyber-security, for the exchange of information and coordination of customer assets must be developed and implemented. Lastly, the customer must be willing to participate with proper incentive.

B. Aggregator

The aggregator is a for-profit market entity engaged in interacting with the customer and the bulk power market in a fully deregulated market structure. As shown in Fig. 1, the aggregator is situated between the DRX and the customer (in a fully deregulated market structure). Note that the DRX can exist in conjunction with existing deregulated market structures. The aggregator energy management system interacts with each of the customer HEMSs. In this paper, we are only considering one aggregator entity, but it is expected that several aggregators may exist within the same distribution area. The existence of an aggregator would depend on legislative policies, but this is beyond the scope of this paper.

The aggregator coordinates a set of participating customers, each with a set of schedulable assets. In this paper, we are currently only considering schedulable loads in the form of smart appliances, but this approach could be extended to other types of assets such as DG, thermal loads (e.g., electric water heaters [25] and heating-ventilation-air-conditioning systems [26]), and electric vehicles (in the form of vehicle-to-grid [27] or scheduling vehicle charging cycles [13]).

The scheduling problem is proposed as a day-ahead optimization. To make decisions, the aggregator requires information about the customer loads, the forecast utility pricing, and the forecast spot market pricing in the bulk electricity market. Using this information, the aggregator must find the CIP and schedule of loads to maximize its profit. Because it is a day-ahead optimization, there are constraints on the execution time of the optimization technique used. This time constraint and the complexity of the scheduling problem (i.e., the class of problems is, in general, NP-complete) due to the large number of customer assets leads to the use of heuristics. Other objectives could be considered, such as minimizing the peak load, or considering multiple objectives in the form of a multi-objective optimization using Pareto-fronts [28]. In this paper, we solely optimize for the aggregator profit to demonstrate that a purely economic motivation will affect the desired change of reduced peak demand on the entire system.

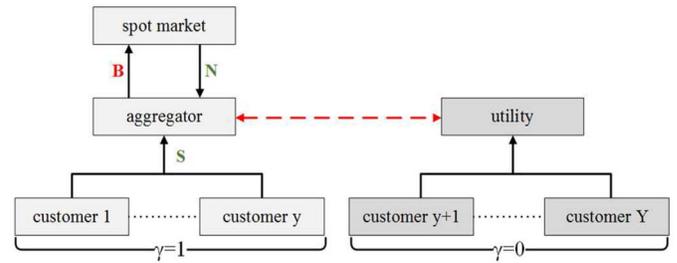


Fig. 2. Money flow with respect to the aggregator, customer, spot market, and utility. The customer has a choice of electricity provider. Customers $\{1 \dots y\}$ pay the CIP to the aggregator for their schedulable loads. Customers $\{y+1 \dots Y\}$ decide the CIP is not worth the inconvenience and purchase electricity from the utility company. The solid arrows represent the money flowing in the system. The dashed red arrow indicates the possible need for a relationship between the aggregator and utility company, which is beyond the scope of this paper.

CIP is a proposed pricing structure that the aggregator would offer all customers to allow the rescheduling of their loads. That is, instead of paying the utility company, the customer pays the aggregator the CIP for electricity. The customer paying the CIP for electricity to the aggregator at the time the asset has been rescheduled to is one part of the profit of the aggregator. The sum of these payments over all customers and all rescheduling events is denoted S . The other two components to the aggregator profit are: 1) the aggregator selling a negative load to the spot market where the assets have been rescheduled from (denoted N); and 2) the aggregator buying spot market electricity where the assets have been rescheduled to (denoted B). This exchange is outlined in Fig. 2. The aggregator would, perhaps, need to enter into a leasing agreement with the utility company for the use of the distribution assets, but modeling this and other potential fixed costs are beyond the scope of this paper.

C. Customer

Each customer under agreement with the aggregator has a set of schedulable loads. In this paper, we are only considering flexible, noninterruptible smart appliances according to the definitions given in [29]. Each customer load has an availability window associated with it. The availability window describes the times during the day that a customer will allow their schedulable load to be rescheduled. In addition to the availability window, each customer has a pricing point that must be met on each load to allow it to be rescheduled. That is, if the price reduction to be received from a rescheduled load at the given CIP is not worth the inconvenience to the customer, the customer may choose to not have their load rescheduled and instead pay the utility company for electricity for that load. This is shown in Fig. 2 as the set of customers $(y+1)$ to Y interacting with the utility instead of the aggregator. Only those loads that are agreed for DR between the customer and aggregator utilize the CIP. The base load and those loads not agreed upon will utilize the status quo of the utility company, e.g., real-time price and time-of-use. This choice of supplier is a powerful new tool for the customer and offers the customer an avenue to participate in the spot

market (through the aggregator entity), which may reduce the customer electricity bill and offer freedom of choice.

III. HEURISTIC FRAMEWORK

A. Overview

We solve the SGRA problem using a heuristic optimization framework, borrowed from concepts in resource allocation in computing, that finds near-optimal solutions to problems. We use heuristic optimization methods because, in general, the class of problems is NP-complete. In this paper, the heuristic framework is designed to be a day-ahead optimization, using a resolution of 15-min intervals. This implies that a given heuristic would need to have a runtime of less than 24-h to be useful. This also gives each vector 96 entries (96 15-min intervals for a complete 24-h period).

B. Schedulable Loads

To reschedule load, the aggregator requires information on the set of schedulable loads. These schedulable loads represent a subset of the system load. For each schedulable load i , the aggregator receives information from the customer on the following:

- 1) δ_i , the runtime duration (in 15-min intervals);
- 2) p_i , the average power rating (in kW);
- 3) t_{i_start} , the customer scheduled start time;
- 4) (A_{i_start}, A_{i_dur}) , a two-tuple that represents the customer-defined availability window for load i determined by the availability window start time, A_{i_start} , and the availability window duration, A_{i_dur} .

In this paper, we assume that the aggregator knows the exact time a load will run (i.e., from t_{i_start} for δ_i time intervals) if it is not rescheduled by the aggregator (i.e., the start time is deterministic). In our future work, this will be represented as a probability distribution based upon historical runtimes (i.e., the start time is stochastic), leading to a stochastic SGRA problem.

C. Aggregator

Let λ be the CIP vector containing 96-elements, where each element λ_t gives the aggregator determined CIP at time interval t . In addition to the information about the schedulable loads, the aggregator possesses information on the following:

- 1) $\gamma(i, \lambda, t)$, a binary function that represents whether the customer will allow load i to be rescheduled to time t with CIP λ ($\gamma = 1$) or not ($\gamma = 0$);
- 2) $s(t)$, the forecast spot market price of electricity in the bulk electricity market (in cents/kWh);
- 3) $r(t)$, the forecast price of electricity from the utility company (in cents/kWh).

Because the customer also has access to the forecast utility price (e.g., real-time price and time-of-use), if the CIP, λ , does not offer enough of a reduction in pricing to justify the inconvenience of rescheduling the load, the customer has the opportunity to refrain from participation, as represented by the binary function, γ . Therefore, the position of the aggregator is to find the following:

- 1) L , the set of loads the aggregator is rescheduling;
- 2) t_{i_resch} , the rescheduled start time for load i ;
- 3) λ , the CIP vector.

So as to maximize profit, given in the following subsection. Let I be the total number of schedulable loads. The cardinality of L is less than or equal to I (i.e., $|L| \leq I$) because the aggregator has information about all I schedulable customer loads, but it does not necessarily have to reschedule all loads.

D. Objective Function

The monetary exchange, representing the aggregator profit, is shown in Fig. 2. For the aggregator, let S be the total income received for selling electricity to customers, given by (1), N be the total income received for selling negative load to the spot market given by (2), and B be the total cost paid to the spot market for buying electricity given by (3). The exact payment received from N would depend on policy, such as the outcome of FERC Order 745 [30] and its future iterations; however, we are not addressing energy policy in this paper. We assume that the aggregator is a well-behaved agent that does not manipulate the market (such as by misrepresenting the sum of the negative load) and is paid the difference from a deterministic baseline load. The calculations for S , N , and B are given as

$$S = \sum_{i \in L} \left[\gamma(i, \lambda, t_{i_resch}) \sum_{t=t_{i_resch}}^{t_{i_resch} + \delta_i - 1} \frac{\lambda_t p_i}{4} \right] \quad (1)$$

$$N = \sum_{i \in L} \left[\gamma(i, \lambda, t_{i_resch}) \sum_{t=t_{i_start}}^{t_{i_start} + \delta_i - 1} \frac{s(t) p_i}{4} \right] \quad (2)$$

$$B = \sum_{i \in L} \left[\gamma(i, \lambda, t_{i_resch}) \sum_{t=t_{i_resch}}^{t_{i_resch} + \delta_i - 1} \frac{s(t) p_i}{4} \right]. \quad (3)$$

The forecast aggregator profit P is given as

$$P = N + S - B. \quad (4)$$

The heuristic optimization problem is set up as follows:

$$\max_{t_{i_resch} \forall i \in L, \lambda = (\lambda_1, \dots, \lambda_{96})} P \quad (5)$$

subject to

$$A_{i_start} \leq t_{i_resch} \leq A_{i_start} + A_{i_dur} \quad \forall i \in L \quad (6)$$

and

$$t_{i_resch} \in \mathbb{Z} \quad \forall i \in L \quad (7)$$

$$\lambda_t \in \mathbb{R} \quad t = 1, \dots, 96. \quad (8)$$

IV. SIMULATION SETUP

A. Overview

The following section describes parameters and models that are used to conduct the simulation study for analysis. The heuristic framework introduced above can be used with any optimization technique, utility pricing mechanism, customer behavior model, and set of customer smart appliances. Although the results show a profit for the aggregator in the considered distribution system, this does not indicate that an aggregator entity would be profitable in all distribution systems; however, the proposed framework can be used to determine this profitability using relevant data.



Fig. 3. Chromosome structure for the GA. The genes $\lambda_1, \dots, \lambda_{96}$ represent the CIP vector, one element for each 15-min interval in the 24-h period. The genes $t_{1,sch}, \dots, t_{I,sch}$ represent the schedule for the I customer loads that are schedulable.

B. GA

In this research, a Genitor [31] version of GA is used to implement the heuristic framework. We use a GA as an example global search heuristic, but any optimization method can be used with the described framework. GAs have been shown to work well in many optimization problems, such as resource allocation in computing [32]–[34], economic dispatch [35], and unit commitment [36]. If multiple objectives are used, the GA can easily be extended to generate Pareto fronts, e.g., with NSGA-II [28], [37].

The implemented chromosome structure is broken into two parts, each with its own gene type, shown in Fig. 3. The first portion of the chromosome is dedicated to the CIP vector, λ , containing 96 genes representing the price (in cents/kWh) for the corresponding 15-min interval. The second portion of the chromosome represents the schedule of loads, containing one gene for each of the I customer schedulable loads. Let $t_{i,sch}$ be a real value in the interval $[0, 1]$ representing the scheduled start time of load i . To obtain the time interval that each load i is scheduled, the following equation is used: $t_{i,resch} = A_{i,start} + t_{i,sch}A_{i,dur}$. If $t_{i,resch} = t_{i,start}$, then the load is not being rescheduled (i.e., $i \notin L$). The $[0, 1]$ representation of $t_{i,sch}$ is used to avoid violating the customer-defined availability constraints of the loads given in (6).

The Genitor version of the GA has a few defining characteristics. In the initial population, no duplicates are allowed to prevent premature convergence. The Genitor is a steady-state algorithm that maintains a ranked list of chromosomes [in this paper, ranked by (4)], leading to implicit elitism, i.e., between generations, the best solutions are kept. In each generation, two parents are selected using the linear bias function (as defined in [31]) leading to the creation of two new children. The linear bias selection function requires a linear bias parameter that is a real value in the interval $(1, 2]$. A linear bias parameter of 1.5 means the best-ranked solution has a 50% greater chance of being selected than the median solution.

After two chromosomes are selected using the linear bias function, two search operators are applied: 1) crossover; and 2) mutation. The former uses a two-point crossover performed on each of the two portions of the chromosome separately. After the crossover is performed, two new children are created. Within each child, every gene has a probability of mutation that will randomly generate a new value for that gene. These two new children are then evaluated in terms of the objective function [given in (5)], inserted into the sorted population, and the worst two chromosomes are trimmed, leading to a fixed population size. The complete algorithm is shown as pseudocode in Fig. 4.

A parameter sweep was used to determine the best parameters to use for the GA in the scope of this problem.

- 1: initialize population
- 2: order population by (4)
- 3: **repeat**
- 4: select two chromosomes via linear bias
- 5: crossover creating two new chromosomes
- 6: mutation
- 7: insert children chromosomes
- 8: trim the two worst performing chromosomes
- 9: **until** stopping criterion
- 10: **return** best chromosome

Fig. 4. Genitor algorithm.

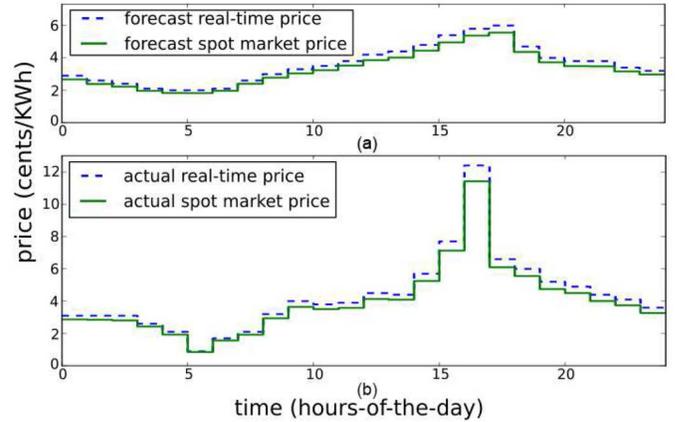


Fig. 5. Real-time [38] and spot market pricing [39] from July 9, 2011. (a) Day-ahead forecast price. (b) Actual price.

The population size was 100, the linear bias parameter was 1.4, and the probability of mutation was 0.01. The stopping criteria was defined as 500 000 total iterations or 10 000 iterations without an increase in the objective function.

Let ω be a real value in the interval $[0, 1]$. To seed the CIP vector, λ , in 50 chromosomes in the initial population, we use the seeding function, denoted $\sigma(t, \omega)$, for each time-window $t = 1, \dots, 96$, given by (9). The schedule for the customer load was randomly generated for each seed. The 50 seeds were generated using values $\omega = n/49, n = 0, \dots, 49$

$$\sigma(t, \omega) = \begin{cases} \omega s(t) & s(t) \geq r(t) \\ \omega r(t) & s(t) < r(t). \end{cases} \quad (9)$$

The rest of the chromosomes in the initial population are randomly generated. For each gene in the CIP vector, representing the cost in cents/kWh at time t , a random value is generated in the interval $[0, \max(r(t), s(t))]$. For each gene in the schedule, representing the scheduled time of load i , a random value in the interval $[0, 1]$ is generated.

C. Pricing Data

The utility pricing and spot market pricing information used in the simulation were real data from Saturday July 9, 2011, obtained from ComEd residential real-time pricing [38] and PJM [39], respectively. This data is given as 24 one-hour intervals. The day-ahead forecast pricing is given in Fig. 5(a) and the actual pricing is given in Fig. 5(b). The data in Fig. 5(a) is used by the GA to determine λ and the schedule of loads.

In Section V, we will evaluate the aggregator profit using the actual price data in Fig. 5(b) with the solution obtained using the forecast price data.

D. Customer

1) *Customer Overview*: In our simulation study, 5555 customers were considered. Each customer has a baseline load and a set of schedulable loads, as described in Section IV-D3. When the aggregator wants to reschedule a customer load, the customer may veto (i.e., $\gamma = 0$) using the process described in the following subsection. In this case study, 56642 loads were available to be rescheduled from the 5555 customers.

2) *Customer Behavior*: A key assumption in the proposed DR methods is customer participation. We model the behavior of each customer for determining whether or not they will allow the aggregator to reschedule their smart appliances using the proposed α -model. In the α -model, each schedulable load i has an associated threshold metric for “customer comfort” in percent, α_i . Let $c_{i,0}$ be the original cost of running load i at the utility real-time price and $c_{i,sch}$ be the rescheduled cost of running load i at the CIP offered by the aggregator. For the owner of load i to allow it to be rescheduled (i.e., $\gamma = 1$), the inequality $c_{i,sch} \leq \alpha_i c_{i,0}$ must hold. This new model allows flexibility for the customer on a load-by-load basis. Additionally, the customer is always guaranteed (if its loads are used by the aggregator), to save $1 - \alpha_i$ times the cost of running load i compared to paying the real-time price. The user inconvenience of the rescheduling of loads is captured through the γ value as opposed to the time dependent models in [40] and [41]. The customer γ values are private, and the aggregator is assumed to operate without receiving this information explicitly.

We use the coefficient-of-variation-based method to generate the α values for each load i , similar to generating task execution times for a heterogeneous suite of machines [42]. We offer an analogous method of generating load α values for a heterogeneous suite of customers. Let μ_a be the desired average load α value for all loads, σ_a be the desired coefficient-of-variation of the load types, and σ_c be the desired coefficient-of-variation of the customers within a load type. For each load type k (given from the rows of Table I), we sample from a Gamma distribution with mean μ_a and standard deviation σ_a to obtain the mean α value for load type k , denoted $\mu_{a,k}$. For each customer that owns load type k , obtain α_i by sampling a Gamma distribution with mean $\mu_{a,k}$ and standard deviation σ_c . This gives similar α values for each type of load, and thus similar customer behavior. This approach was taken because it is assumed that customers will act similar regarding the use of load types (e.g., more flexible with laundry and less flexible with the TV).

A parameter sweep was performed on the input values μ_a , σ_a , and σ_c . A representative result is shown in Section V using the inputs $\mu_a = 0.75$, $\sigma_a = 0.10$, and $\sigma_c = 0.05$. In general, the magnitude of the CIP is sensitive and positively correlated to μ_a (i.e., as μ_a increases, the CIP proportionally increases with respect to the real-time price). Values of σ_a and σ_c are positively correlated with the noise level of the CIP.

TABLE I
SCHEDULABLE SMART APPLIANCES

Penetration (%)	Mean power (kW)	Power std. dev. (kW)	Duration (15-minute intervals)	Start mean (hour)	Start std. dev. (hour)
70	0.5	.05	4	7	1
70	0.5	.05	4	14	3
70	0.5	.05	4	17	1
50	0.75	.1	3	7	1
50	0.75	.1	3	14	3
50	0.75	.1	3	17	1
30	1.0	.2	2	7	1
30	1.0	.2	2	14	3
30	1.0	.2	2	17	1
100	0.25	.01	8	7	1
100	0.25	.01	8	14	3
100	0.25	.01	8	17	1
10	1.5	.3	2	7	1
10	1.5	.3	2	14	3
10	1.5	.3	2	17	1
80	0.4	.05	6	7	1
80	0.4	.05	6	14	3
80	0.4	.05	6	17	1

3) *Customer Loads*: Two types of loads are assumed to be available for each customer in this paper: 1) baseline and 2) schedulable (smart) appliances. The baseline load is divided into thermal, modeled as air conditioning [43] and electric water heaters [25], and other nonschedulable loads. The nonschedulable loads are probabilistically generated for each customer based on the data in [15].

A probabilistic model for 18 generic schedulable appliance types is given in Table I. The penetration level gives the probability that an appliance is present for a given customer; if it is present, the rated power of the appliance, as well as the start hour, is obtained from a normal distribution. Values in Table I were chosen so that the total load reflects actual energy use of an average household. Similar to the nonschedulable loads, a set of schedulable loads corresponding to each customer is generated probabilistically using the data in Table I.

Each probabilistically generated load i has an associated availability window, $(A_{i,start}, A_{i,dur})$, that describes the time-window that the customer has allocated for load i to be scheduled. Recall that $t_{i,start}$ is the originally scheduled starting time for load i . Let $\mathcal{U}(\delta_i, 96)$ be a uniform random variable in the interval $[\delta_i, 96]$. In this paper, to generate the availability window for each load i , an interval of duration $\mathcal{U}(\delta_i, 96)$ is generated around the starting time $t_{i,start}$. That is, $A_{i,dur} = \mathcal{U}(\delta_i, 96)$ and $A_{i,start} = t_{i,start} - A_{i,dur}/2$.

V. RESULTS

A total of 56642 schedulable loads (i.e., $I = 56642$) from the 5555 customers were randomly generated using the data from Table I. The schedulable customer loads correspond to 11.2% of the total energy used by the 5555 customers. To capture the algorithm in steady state, a 2-h window was added to the start and end of the simulation. Any appliance load that occurs within these windows did not contribute toward the objective function (i.e., only the 24-h window was used for the objective function calculation). The GA ran for 375 000 iterations before terminating, taking 113 min on an Intel i7 4900MQ processor running at 2.8 GHz using

a C++ implementation in Ubuntu Linux. The final objective value, i.e., forecast aggregator profit, was $P = \$813.92$ [based on Fig. 5(a)]. When evaluated for the actual real-time and spot market pricing, the schedule determined by the GA resulted in an aggregator profit of $\$947.90$ [based on Fig. 5(b)]. This increase in profit from forecast to actual is because the actual spot market pricing at the peak period was much larger than forecast (as shown in Fig. 5), leading to an increase in profit from the N component of the profit function. From a customer standpoint, the total savings of all 5555 customers was $\$460.31$ and $\$794.93$ when using the forecast and actual data, respectively, for the 24-h period under consideration. The increase in savings is also due to the large increase in peak real-time price that the customer no longer has to pay. For the settlement of the customer DR, the aggregator uses the actual spot market price data. The total customer savings implies an average saving of $\$0.14$ per customer with a range of savings between $\$0.02$ and $\$0.33$. This range is indicative of the possible monetary benefits from the customer being more flexible with their loads (in the availability of the load and the customer α values) and bringing more energy (i.e., a greater number of assets) to the aggregator to participate in DR. Although the average daily savings may appear small, in this paper, we are focusing on the viability of the aggregator.¹ In general, the aggregator makes less profit and the customer saves more when the μ_a value is decreased, and vice-versa.

Fig. 6 shows the change in the load before and after the optimization occurs. Fig. 6(a) compares the system load before and after the optimization. As hypothesized, if the aggregator entity optimizes purely for economic reasons, the overall change in the system peak load may be beneficial, as is the case in this paper. The aggregator-based residential DR program was able to reduce the peak of the participating 5555 customers by 12.5%, resulting in a 2.66 MW reduction at 4:45 P.M. In Fig. 6(b), the portion of the load that is schedulable is shown. The area under the curve is 11.2% of the total system energy, with 19.4% of the total load reschedulable at the peak. This figure shows in greater resolution when the rescheduling of customer loads occurs. Over half of the schedulable load at the peak is moved off-peak. The reason this value is not higher is due to the customer availability windows described in Section III-B. Because of this constraint, not all of the load can be moved to off-peak hours. Fig. 6(c) explicitly shows the difference in load between the system before and after the DR. The green shaded regions with the “/” hashing are the areas that the load was reduced, corresponding to the reduction in the peak. In the other areas, shaded red with “\” hashing, the load was increased, corresponding to the load moving to off-peak hours. The large negative difference in load directly corresponds to the component of aggregator profit obtained by selling negative load, N , to the spot market. The positive difference in load is the portion of the load that contributes to the $S - B$ component of the aggregator profit function.

¹Monetary benefits, however, are not the only reason that early adopters may want to participate. It has been shown that often customers are motivated by altruistic reasons, such as environmental benefits (i.e., “being green”) [44].

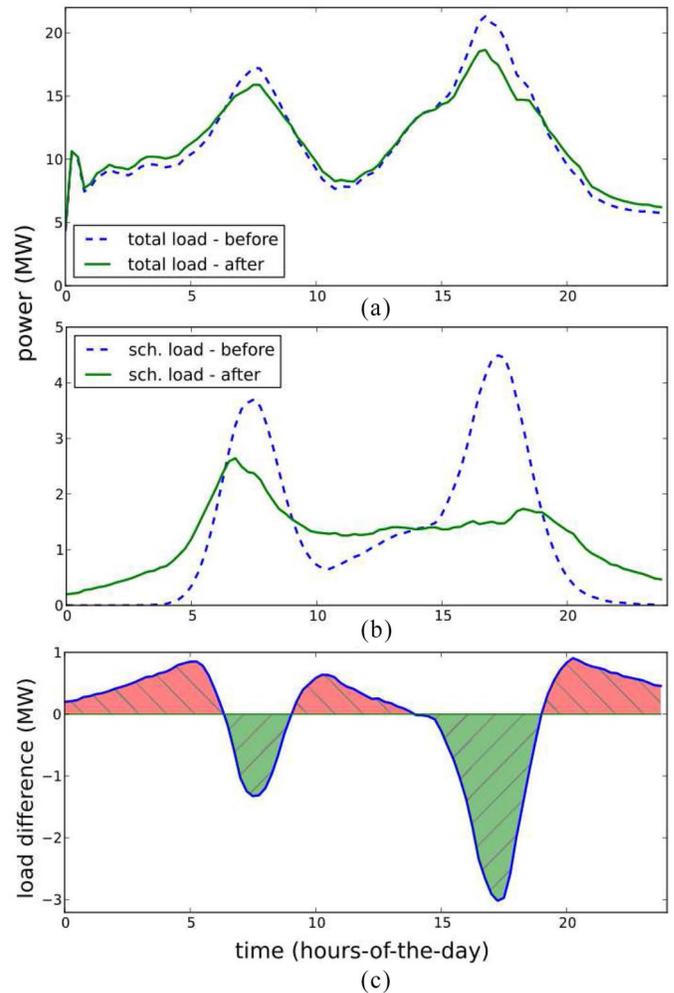


Fig. 6. Change in load from before and after the aggregator DR action. (a) Overall system load of the 5555 customers. (b) Schedulable load. (c) Difference in load (i.e., after minus before).

The customer pricing incentive obtained by the optimization, is shown in Fig. 7. Fig. 7(a) shows the incentive pricing compared to the forecast real-time and spot market prices. The CIP is lower than the forecast real-time price and, in general, the actual real-time price [Fig. 7(b)]. This indicates the customer receives a competitive, and reduced, rate of electricity for participating with the aggregator [including a hedge against the risk of large price spikes in the real-time price, such as at 4 P.M. in Fig. 7(b)].

To estimate the response of the spot market to the aggregator DR, a pseudo-spot-market-response is emulated. A sixth-order polynomial was fit to the PJM forecast spot market price for July 9, 2011 describing the spot market price (in cents/kWh) as a function of the load (in MW). The coefficients of the polynomial are given in the appendix. The adjusted R^2 value for the polynomial to the data is 0.987, indicating a close fit. The l^2 -norm of the difference between the price determined by the polynomial and the original forecast price is 0.516 cents/kWh. Note that this is a simple model for quantifying the changes the DR has on the market for July 9, 2011, and should not be generalized as a predictive tool.

The difference in system load due to the DR action [i.e., Fig. 6(c)] was added to the forecast PJM clearing

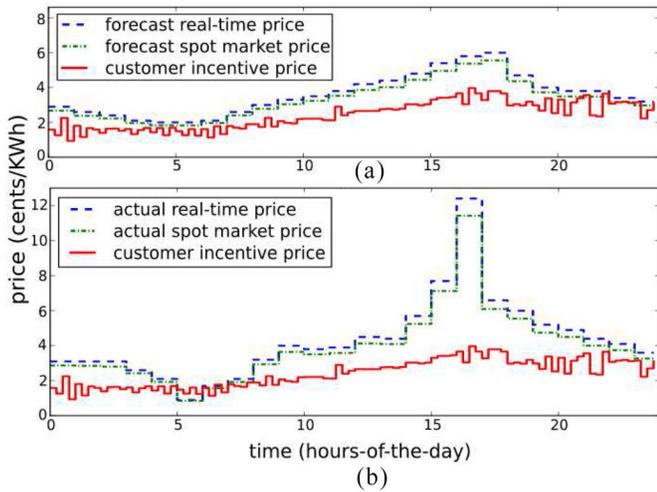


Fig. 7. Real-time and spot market pricing compared to the CIP. CIP compared to the (a) day-ahead forecast price and (b) actual price.

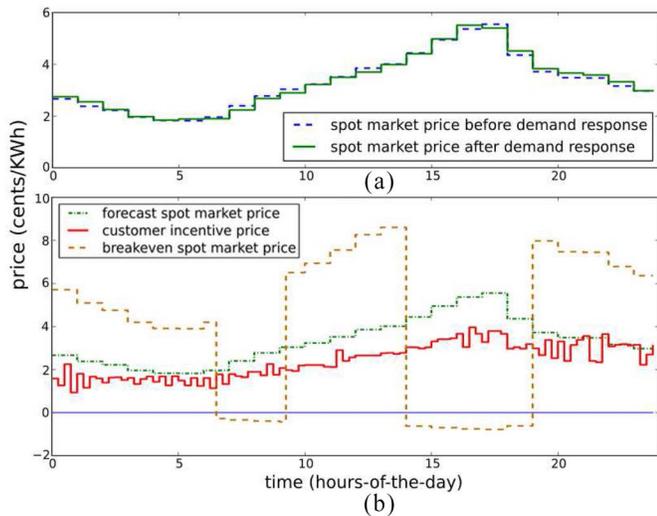


Fig. 8. Results of the pseudo-market-response derived from the sixth-order polynomial regression model. (a) Predicted change in forecast spot market price as a result of the DR. (b) Spot market price that will make aggregator break-even (no profit).

load (available in [39]) and the polynomial fit was used to determine the resultant spot market price. The change in spot market price due to the DR is visualized in Fig. 8(a). The l^2 -norm of the difference between the emulated pseudo-spot-market-response and the forecast spot market price is 0.517 cents/kWh. When compared to the original l^2 -norm of 0.516, this indicates the change in spot market price from the DR from one aggregator entity is small. However, when many aggregators exist within the purview of a single ISO, this market response will need to be investigated further.

To determine the breakeven point for the profit of the aggregator with the given DR, the forecast price was scaled until $P = \$0$. A scalar value, β , was applied in a positive manner at the times the load was increased and a negative manner at the times the load was decreased, the red and green areas in Fig. 6(c), respectively. This was done to increase the cost from the B term in the red shaded areas and to decrease the profit

from the N term in the green shaded areas. The breakeven price was determined with $\beta = 1.141$, given in Fig. 8(b).

VI. CONCLUSION

We proposed an aggregator-based residential DR approach for scheduling residential customer assets. A CIP structure was proposed to compensate the customer for the inconvenience of rescheduling their assets. This new pricing structure gives the customer a near real-time choice of electricity supplier in a fully deregulated market scenario. A heuristic framework was designed to perform an optimization on the profit of the aggregator. To validate the heuristic framework, a system comprised of 5555 customer households and 56 498 schedulable loads was simulated using a GA implementation of the framework. The CIP found by the GA was, in general, lower at all times than the customer would pay via real-time pricing. Despite this, the aggregator was able to make a profit by selling negative peak load to the spot market. This showed an example of optimizing for purely economical reasons in the form of aggregator profit, and enacting an overall change on the system peak load. This change benefits the customer of the aggregator (in the form of reduced cost of electricity for schedulable loads), the aggregator (in the form of a profit), and also those customers not participating with the aggregator (because the overall system peak is lowered as a common good).

APPENDIX

MARKET RESPONSE APPROXIMATION

Let $s_p(x)$ be the predicted spot market price in cents/kWh for a given load x in MW. The sixth-order polynomial approximation of $s_p(x)$ for ComEd within PJM for July 9, 2011, is given in (10). Note that the regression fit is empirical over the domain $x = [69\ 200, 112\ 000]$ MW and should not be used for loads outside of these values

$$s_p(x) = 1.84 \times 10^{-26}x^6 - 9.60 \times 10^{-21}x^5 + 2.07 \times 10^{-15}x^4 - 2.38 \times 10^{-10}x^3 + 1.52 \times 10^{-5}x^2 - 0.516x + 7230 \quad (10)$$

for $69\ 200 \leq x \leq 112\ 000$.

ACKNOWLEDGMENT

The authors would like to thank M. Oxley and E. Jonardi for their valuable comments.

REFERENCES

- [1] Litos Strategic Communication, *The Smart Grid: An Introduction*, U.S. Dept. Energy, Washington, DC, USA, Tech. Rep. 560.01.04, 2009. [Online]. Available: <http://energy.gov/oe/downloads/smart-grid-introduction-0>
- [2] J. J. Conti, P. D. Holtberg, J. A. Beamon, S. A. Napolitano, A. Michael Schaal, and J. T. Turnure, *Annual Energy Outlook 2013*, U.S. Energy Inf. Admin., Tech. Rep. AEO2013, Apr. 2013. [Online]. Available: [http://www.eia.gov/forecasts/archive/aeo13/pdf/0383\(2013\).pdf](http://www.eia.gov/forecasts/archive/aeo13/pdf/0383(2013).pdf)
- [3] J. Osborne and D. Warrier, "A Primer on Demand Response—The Power Grid: Evolving From a Dumb Network to a Smart Grid," Thomas Weisel Partners Equity Research, San Francisco, CA, USA, Oct. 2007. [Online]. Available: http://downloads.lightreading.com/internetevolution/Thomas_Weisel_Demand_Response.pdf

- [4] PJM Forward Market Operations. (Jun. 2014). *Energy and Ancillary Services Market Operations*. [Online]. Available: <http://www.pjm.com/~media/documents/manuals/m11.ashx>
- [5] P. C. Stern, "Information, incentives, and proenvironmental consumer behavior," *J. Consum. Policy*, vol. 22, no. 4, pp. 461–478, Dec. 1999.
- [6] E. G. Coffman, *Computer and Job-Shop Scheduling Theory*. Hoboken, NJ, USA: Wiley, 1976.
- [7] D. Fernandez-Baca, "Allocating modules to processors in a distributed system," *IEEE Trans. Softw. Eng.*, vol. 15, no. 11, pp. 1427–1436, Nov. 1989.
- [8] O. Ibarra and C. E. Kim, "Heuristic algorithms for scheduling independent tasks on nonidentical processors," *J. ACM*, vol. 24, no. 2, pp. 280–289, Apr. 1977.
- [9] X. Chen, T. Wei, and S. Hu, "Uncertainty-aware household appliance scheduling considering dynamic electricity pricing in smart home," *IEEE Trans. Smart Grid*, vol. 4, no. 2, pp. 932–941, Mar. 2013.
- [10] A. Agnetis, G. de Pascale, P. Detti, and A. Vicino, "Load scheduling for household energy consumption optimization," *IEEE Trans. Smart Grid*, vol. 4, no. 4, pp. 2364–2373, Dec. 2013.
- [11] T. H. Chang, M. Alizadeh, and A. Scaglione, "Coordinated home energy management for real-time power balancing," in *Proc. IEEE Power Energy Soc. Gen. Meeting*, San Diego, CA, USA, 2012, pp. 1–8.
- [12] C. Chen, J. Wang, Y. Heo, and S. Kishore, "MPC-based appliance scheduling for residential building energy management controller," *IEEE Trans. Smart Grid*, vol. 4, no. 3, pp. 1401–1410, Aug. 2013.
- [13] M. A. A. Pedrasa, T. D. Spooner, and I. F. MacGill, "Coordinated scheduling of residential distributed energy resources to optimize smart home energy services," *IEEE Trans. Smart Grid*, vol. 1, no. 2, pp. 134–143, Sep. 2010.
- [14] T. Logenthiran, D. Srinivasan, and T. Z. Shun, "Demand side management in smart grid using heuristic optimization," *IEEE Trans. Smart Grid*, vol. 3, no. 3, pp. 1244–1252, Sep. 2012.
- [15] R. Roche, *Agent-based Architectures and Algorithms for Energy Management in Smart Grids*, Ph.D. dissertation, Sciences Pour l'Ingénieur et Microtechniques, Institut de Recherche sur les Transports, l'Énergie et la Société—Laboratoire Systèmes et Transports, Univ. Technol. Belfort-Montbéliard, Belfort, France, Dec. 2012.
- [16] M. Roozbehani, M. A. Dahleh, and S. K. Mitter, "Volatility of power grids under real-time pricing," *IEEE Trans. Power Syst.*, vol. 27, no. 4, pp. 1926–1940, Nov. 2012.
- [17] C. Quinn, D. Zimmerle, and T. H. Bradley, "The effect of communication architecture on the availability, reliability, and economics of plug-in hybrid electric vehicle-to-grid ancillary services," *J. Power Sources*, vol. 195, no. 5, pp. 1500–1509, Mar. 2010.
- [18] C. Quinn, D. Zimmerle, and T. H. Bradley, "An evaluation of state-of-charge limitations and actuation signal energy content on plug-in hybrid electric vehicle, vehicle-to-grid reliability, and economics," *IEEE Trans. Smart Grid*, vol. 3, no. 1, pp. 483–491, Mar. 2012.
- [19] L. Gkatzikis, I. Koutsopoulos, and T. Salonidis, "The role of aggregators in smart grid demand response markets," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 7, pp. 1247–1257, Jul. 2013.
- [20] Z. Yu, L. Jia, M. C. Murphy-Hoye, A. Pratt, and L. Tong, "Modeling and stochastic control for home energy management," *IEEE Trans. Smart Grid*, vol. 4, no. 4, pp. 2244–2255, Dec. 2013.
- [21] P. Koponen, J. Ikäheimo, A. Vicino, A. Agnetis, G. De Pascale, N. R. Carames, J. Joseba, E. F. Sánchez-Úbeda, P. Garcia-Gonzalez, and R. Cossent, "Toolbox for aggregator of flexible demand," in *Proc. Future Energy Grids Syst. Symp.*, Florence, Italy, Sep. 2012, pp. 623–628.
- [22] R. Belhomme, R. Cerero, G. Valtorta, and P. Eyrolles, "The ADDRESS project: Developing active demand in smart power systems integrating renewables," in *Proc. IEEE PES Gen. Meeting*, San Diego, CA, USA, Jul. 2011, pp. 1–8.
- [23] California Energy Commission. (Mar. 28, 2014). *2013 Integrated Energy Policy Report*. [Online]. Available: http://www.energy.ca.gov/2013_energypolicy/
- [24] T. Hansen, R. Roche, S. Suryanarayanan, H. J. Siegel, D. Zimmerle, P. M. Young, and A. A. Maciejewski, "A proposed framework for heuristic approaches to resource allocation in the emerging smart grid," in *Proc. IEEE PES Int. Conf. Power Syst. Technol. (POWERCON)*, Auckland, New Zealand, Oct. 2012, pp. 1–6.
- [25] P. Du and N. Lu, "Appliance commitment for household load scheduling," *IEEE Trans. Smart Grid*, vol. 2, no. 2, pp. 411–419, Jun. 2011.
- [26] K. F. Fong, V. I. Hanby, and T. T. Chow, "HVAC system optimization for energy management by evolutionary programming," *Energy Build.*, vol. 38, no. 3, pp. 220–231, Mar. 2006.
- [27] E. Sortomme and M. A. El-Sharkawi, "Optimal scheduling of vehicle-to-grid energy and ancillary services," *IEEE Trans. Smart Grid*, vol. 3, no. 1, pp. 351–359, Mar. 2012.
- [28] R. Friese, T. Brinks, C. Oliver, A. A. Maciejewski, and H. J. Siegel, "Analyzing the trade-offs between minimizing makespan and minimizing energy consumption in a heterogeneous resource allocation problem," in *Proc. 2nd Int. Conf. Adv. Commun. Comput.*, Oct. 2012, pp. 81–89.
- [29] Z. Chen, L. Wu, and Y. Fu, "Real-time price-based demand response management for residential appliances via stochastic optimization and robust optimization," *IEEE Trans. Smart Grid*, vol. 3, no. 4, pp. 1822–1831, Dec. 2012.
- [30] D. Hunger and D. Hough, *FERC Order 745: Demand Response Compensation in Organized Wholesale Energy Markets*. U.S. Fed. Energy Regul. Commiss., Tech. Rep. RM10-17-000, Mar. 2011.
- [31] D. Whitley, "The GENITOR algorithm and selective pressure: Why rank-based allocation of reproductive trials is best," in *Proc. 3rd Int. Conf. Genet. Algorithms*, San Francisco, CA, USA, Jun. 1989, pp. 116–121.
- [32] L. D. Briceno, H. J. Siegel, A. A. Maciejewski, M. Oltikar, J. Brateman, J. White, J. Martin, and K. Knapp, "Heuristics for robust resource allocation of satellite weather data processing on a heterogeneous parallel system," *IEEE Trans. Parallel Distrib. Syst.*, vol. 22, no. 11, pp. 1780–1787, Nov. 2011.
- [33] T. D. Braun, H. J. Siegel, N. Beck, L. Boloni, M. Maheswaran, A. I. Reuther, J. P. Robertson, M. D. Theys, B. Yao, D. Hengsen, and R. F. Freund, "A comparison of eleven static heuristics for mapping a class of independent tasks onto heterogeneous distributed computing systems," *J. Parallel Distrib. Comput.*, vol. 61, no. 6, pp. 810–837, Jun. 2001.
- [34] F. D. Croce, R. Tadei, and G. Volta, "A genetic algorithm for the job shop problem," *Comput. Oper. Res.*, vol. 22, no. 1, pp. 15–24, Jan. 1995.
- [35] D. C. Walters and G. B. Sheble, "Genetic algorithm solution of economic dispatch with valve point loading," *IEEE Trans. Power Syst.*, vol. 8, no. 3, pp. 1325–1332, Aug. 1993.
- [36] S. A. Kazarlis, A. G. Bakirtzis, and V. Petridis, "A genetic algorithm solution to the unit commitment problem," *IEEE Trans. Power Syst.*, vol. 11, no. 1, pp. 83–92, Feb. 1996.
- [37] D. Kalyanmoy, A. Pratap, S. Agarwal, and T. A. M. T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Trans. Evol. Comput.*, vol. 6, no. 2, pp. 182–197, Apr. 2002.
- [38] ComEd. (Aug. 18, 2013). *ComEd Residential Real-Time Pricing Program*. [Online]. Available: <https://rtrp.comed.com/live-prices/>
- [39] PJM Daily Report Archives. (Jan. 14, 2014). *U.S. Federal Energy Regulatory Commission*. [Online]. Available: <http://www.ferc.gov/market-oversight/mkt-electric/pjm/pjm-iso-archives.asp>
- [40] L. Gkatzikis, T. Salonidis, N. Hegde, and L. Massoulié, "Electricity markets meet the home through demand response," in *Proc. IEEE Conf. Decis. Control*, Maui, HI, USA, Dec. 2012, pp. 5846–5851.
- [41] T. Bapat, N. Sengupta, S. K. Ghai, V. Arya, Y. B. Shrinivasan, and D. Seetharam, "User-sensitive scheduling of home appliances," in *Proc. ACM SIGCOMM Workshop Green Netw.*, Toronto, ON, Canada, Aug. 2011, pp. 43–48.
- [42] S. Ali, H. J. Siegel, M. Maheswaran, D. Hengsen, and S. Ali, "Representing task and machine heterogeneities for heterogeneous computing systems," *Tamkang J. Sci. Eng.*, vol. 3, no. 3, pp. 195–208, Nov. 2000.
- [43] N. Lu and Y. Zhang, "Design considerations of a centralized load controller using thermostatically controlled appliances for continuous regulation reserves," *IEEE Trans. Smart Grid*, vol. 4, no. 2, pp. 914–921, Jun. 2013.
- [44] A. Zipperer, P. A. Aloise-Young, S. Suryanarayanan, R. Roche, L. Earle, D. Christensen, P. Bauleo, and D. Zimmerle, "Electric energy management in the smart home: Perspectives on enabling technologies and consumer behavior," *Proc. IEEE*, vol. 101, no. 11, pp. 2397–2408, Aug. 2013.



Timothy M. Hansen (GSM'11) received the B.S. degree in computer engineering from the Milwaukee School of Engineering, Milwaukee, WI, USA, in 2011. He is currently pursuing the Ph.D. degree from the Department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO, USA.

He is with the Robust Computing and Advanced Power Engineering Laboratory groups, Colorado State University. His current research interests

include robust resource allocation and optimization applied to cyber-physical-social systems, power engineering, and high-performance computing.



Robin Roche (GSM'09–M'13) received the Ph.D. degree in electrical engineering from the Université de Technologie de Belfort-Montbéliard (UTBM), Belfort, France, in 2012.

He is currently an Associate Professor with the Research Institute on Transportation, Energy, and Society, and with the Department of Energy, UTBM. His current research interests include energy management in power and energy systems.



Siddharth Suryanarayanan (S'00–M'04–SM'10) received the Ph.D. degree in electrical engineering, Arizona State University, Tempe, AZ, USA, in 2004.

He is an Associate Professor with the Department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO, USA, where he teaches and performs research in the area of advanced electric power systems.



Anthony A. Maciejewski (F'05) received the B.S.E.E., M.S., and Ph.D. degrees in electrical engineering from Ohio State University, Columbus, OH, USA, in 1982, 1984, and 1987, respectively.

From 1988 to 2001, he was a Professor of Electrical and Computer Engineering at Purdue University, West Lafayette, IN, USA. He is currently a Professor and the Department Head of Electrical and Computer Engineering at Colorado State University, Fort Collins, CO, USA. His current research interests include robotics and high performance computing.



Howard Jay (“H. J.”) Siegel (F'90) received the double B.S. degrees in electrical engineering and management from the Massachusetts Institute of Technology, Cambridge, MA, USA, and the M.A., M.S.E., and Ph.D. degrees in electrical engineering and computer science from Princeton University, Princeton, NJ, USA, in 1972, 1974, and 1977, respectively.

Since 2001, he has been the Abell Endowed Chair Distinguished Professor of Electrical and Computer Engineering at Colorado State University, Fort Collins, CO, USA, where he is a Professor of Computer Science. From 1976 to 2001, he was a Professor at Purdue University, West Lafayette, IN, USA. His current research interests include robust energy-aware resource management for heterogeneous parallel and distributed computing, the smart grid for electricity, parallel algorithms, parallel machine interconnection networks, and reconfigurable parallel computer systems. He has co-authored over 420 technical papers. He was a consultant for industry and government.

Prof. Siegel was the Co-Editor-in-Chief of the *Journal of Parallel and Distributed Computing*. He was on the Editorial Boards of the IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS and the IEEE TRANSACTIONS ON COMPUTERS. He is a Fellow of the Association for Computing Machinery.