

iEMS for Large Scale Charging of Electric Vehicles: Architecture and Optimal Online Scheduling

Shiyao Chen, and Lang Tong
 School of Electrical and Computer Engineering
 Cornell University, Ithaca, NY 14853
 Email: {sc933,lt35}@cornell.edu

Abstract—The problem of large scale charging of electric vehicles (EVs) is considered. An architecture for the energy management system (EMS) is proposed based on the concept of network switched charging where chargers are controlled by a scheduler that optimizes the overall operating profit of the service provider. It is assumed that the EMS has access to collocated renewable sources (e.g. solar power) and can supplement the renewable with purchased electricity from the grid. The renewable source may vary arbitrarily, and requests of all EVs accepted for service must be completed by their respective deadlines.

Under a deterministic model for arbitrary arrivals, charging requests, and service deadlines, online scheduling of EV charging is formulated as a multi-processor deadline scheduling problem for which the optimal scheduler maximizes the competitive ratio against the best offline scheduler. An online scheduling algorithm, referred to as TAGS, is proposed based on the principle of threshold admission and greedy scheduling. TAGS has the complexity of $O(n \log n)$ where n is the number of EVs in the facility. It is shown that, when the price offered to the EV customers is higher than the purchasing price of electricity from the grid, TAGS achieves the competitive ratio of 1. Otherwise, TAGS achieves the maximum competitive ratio given by the inverse of a real root of a certain polynomial. Simulations are used to evaluate the performance of TAGS against standard benchmarks and for the setting of optimal charging price.

Index Terms—Charging of electric vehicles; energy management systems; deadline scheduling; competitive ratio analysis; renewable integration.

I. INTRODUCTION

WE consider the problem of charging of electric vehicles (EVs) at parking facilities where the charging of a large number of EVs can be centrally managed. We assume that EVs arrive with arbitrary charging requests, and each request has to be completed by a certain deadline. The charging service provider (CSP) meets demands using collocated renewable sources (e.g., solar) supplemented by electricity purchased from the grid.

The CSP aims to maximize the operating profit. Achieving this objective, however, involves complex decisions with interacting parameters. It is thus necessary to develop an intelligent energy management system (iEMS) that procures energy from different sources, dispatches of energy to EVs, makes admission decisions on new arrivals, and determines the price of charging.

Research was sponsored by the National Science Foundation under a Grant CNS-1135844.

A. An iEMS Architecture

We propose an iEMS architecture illustrated in Figure 1 that has some of the key functions discussed above. The key characteristic of the proposed architecture is the *network switched charging* where a scheduler controls a switch that optimizes charging schedules based on system capacity, and available renewable sources and local storage.

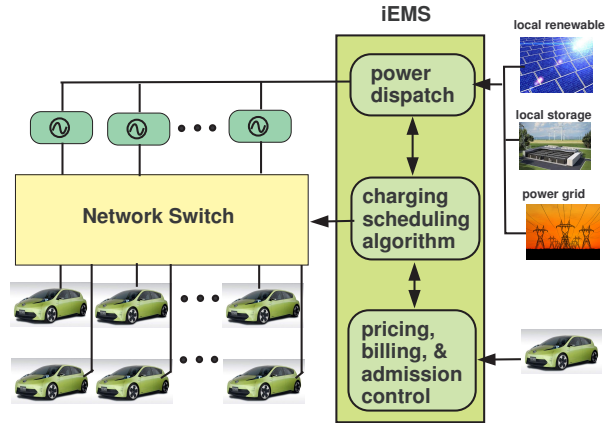


Fig. 1. Architecture for network switched charging and iEMS.

The objective of the iEMS is maximizing the operating profit for the CSP. The hardware system of the proposed iEMS includes a dispatcher that delivers power from a mix of sources—renewable, purchased electricity, and possibly local storage—to a set of chargers. The iEMS also includes a switch that connects chargers to EVs admitted to the facility. Through the network switch, a scheduler activates EV charging with preemption, making it possible for the iEMS to serve more urgent or more profitable requests.

The iEMS is run by its software system that makes engineering and economic decisions. At the core of the software system for the iEMS is the charging scheduling algorithm, which is the focus of this paper. The scheduler (i) controls the power dispatcher to procure energy from available sources, (ii) sets the connections of the switch so that a subset of EVs are charged by the available chargers, and (iii) determines the admission of new EVs based on its operating condition. The software system also has to handle billing and other ancillary

services, which are not discussed in this paper.

The optimizations involved in the iEMS are of the mixed integer type. Since the charging of an EV may be interrupted by a more urgent request, deadlines of service requests play a critical role. By specifying deadlines at the time of arrival, the EV customers provide the flexibilities that can be exploited. For example, the scheduler should charge as much as possible when there is ample inexpensive renewable energy. But when electricity has to be purchased to fulfill charging requests, the scheduler can optimize on the time and quantity of purchase. Such an optimization are made by taking into account various factors such as the deadlines of charging requests and possibly the forecast of available renewable in the future.

B. Summary of Results and Contexts

Having described the iEMS architecture, we focus on the scheduling module in Figure 1. In this paper, we adopt a deterministic formulation where EV arrivals, charging demands, charging deadlines, available renewable sources, and the price of electricity from the grid are all deterministic and arbitrary. The decision variables include which EV to charge and by which charger, how long each charging should last, how much electricity needs to be purchased from the grid, and whether a new arrival should be admitted to the system.

We restrict to the design of an *online scheduling algorithm* where the scheduler has the past history but without foresight into future requests. In such a deterministic setting where all scenarios are possible, the approach of *competitive ratio analysis* is appropriate, where the performance of an online algorithm is compared against the optimal offline algorithm for the worst operating scenario.

The competitive ratio (see its definition in Sec II) provides a level of performance guarantee. In particular, if an algorithm has a competitive ratio of 1, it implies that the algorithm is optimal among all online and offline algorithms. If an algorithm achieves the *optimal* competitive ratio $C < 1$, it guarantees better than C fraction of the reward provided by the optimal offline algorithm. And there is no other online algorithm that can be strictly better for all service requests.

We characterize the optimal competitive ratio and present the optimal scheduling algorithm. Assuming that the renewable source can support M chargers, the main result—Theorem 1 in Sec. II—shows that the optimal competitive ratio has two operating regimes. When the price of charging offered to EV customers is higher than that of the purchased electricity, the maximum achievable competitive ratio is one. Otherwise, the maximum achievable competitive ratio is given by the inverse of a real root of a certain polynomial.

The more practically significant result is an online scheduling algorithm that achieves the optimal competitive ratio (see Theorem 2 in Sec. II). The algorithm is referred as TAGS, which is the abbreviation for Threshold Admission and Greedy Scheduling. TAGS performs simple threshold tests on profitability for admission control and makes greedy scheduling decisions. As an online scheduling algorithm, TAGS has the

complexity of $O(n \log n)$ where n is the number of EVs in the facility.

The deterministic model adopted in this paper has the advantage of being agnostic to specific models of traffic, demand, availability of renewable, and battery technology. The competitive ratio analysis based on the worst case service requests, however, is less informative to practitioners; often, it is the average performance that a CSP or an EV customer desires. In this paper we test TAGS using commonly accepted statistical models. While TAGS does not guarantee optimal average performance, it outperforms standard benchmarks such as earliest deadline first (EDF), the least laxity first (LLF), and the unmanaged charging (UC) algorithms.

The profit of the CSP depends on the purchase price of electricity and the offered price to EV customers; their relative values determine the optimal competitive ratio. However, neither can be set arbitrarily by the CSP. In evaluating the performance as a function of pricing, we assume that the CSP is a price taker when it purchases electricity, and EV customers behave with a downward slope demand curve. With such a setting, we evaluate the optimal offered price to the EV customers using simulations.

Due to space limitation, the proof of the main theorem is omitted and is reported in [1].

C. Related Work

There is a growing literature on various aspects of the EV charging problem including the scheduling of chargers; see, e.g., [2, 3, 4, 5, 6, 7]. Most reported work, however, focuses on residential charging scenarios that have somewhat different characteristics from that considered here where EVs are located and managed centrally. The authors of [8] conduct simulation studies of EMS for EV charging at parking garages where heuristic algorithms for scheduling are used. The EV charging for public garages is also considered in [9] where the objective is maximizing the service throughput, and the cost of energy is not part of the optimization.

Subramanian *et al.* consider a setup closest to the present paper in [10] where the scheduling of EV charging is optimized using a combination of renewable energy and energy from the grid. The authors in [10] investigate the EDF, LLF, and a receding horizon control scheduling algorithm; the latter exploits the forecast of renewable energy. Our approach differs from [10] in the use of admission control in scheduling and the objective of the optimization. In particular, [10] is concerned about the peak power drawn from the grid whereas the objective in this paper is maximizing the operating profit.

Prior to this work, large scale charging of electric vehicles is formulated as a deadline scheduling problem with a single processor, and an early version of the single processor TAGS is presented [11]. For variable renewable sources, the generalization of the single processor case is reported in [12, 13]. The present paper strengthens the optimal competitive ratio for the multiprocessor deadline scheduling algorithm.

The technique presented in this paper follows the deadline scheduling framework, originated from the seminal work of

Liu and Layland [14]. The best known on-line scheduling algorithms include the earliest deadline first (EDF) algorithm [14, 15] and the least laxity first (LLF) algorithm [16], both are known to have competitive ratio one when the system is “under-loaded” and there is a single processor. For general operating conditions, EDF and LLF can perform poorly [17], especially when there are overwhelming arrivals. The single processor optimal online scheduling algorithms are proposed by Baruah *et. al.* and Koren and Shasha separately in [18, 19], and $1/4$ is shown to be the maximum competitive ratio. The key difference between the setting of [18, 19] and that considered here is that, in [18, 19], the unfinished jobs have no impact on the performance.

As a deadline scheduling algorithm, our approach presented in this and earlier papers breaks some new ground. The problem considered here corresponds to deadline scheduling involving multiple processors, for which results on competitive ratio are limited and optimal scheduling algorithm unknown. See, *e.g.*, [20]. The main result presented here appears to be the first that, for the multiprocessor deadline scheduling problem, gives a complete characterization of the optimal competitive ratio and an algorithmic construction to achieve the maximum competitive ratio. We obtain these results thanks to including admission control as part of the optimization.

Admission control is natural for the EV charging application where unfinished jobs should be penalized. Consequently, the CSP should not accept all requests, not only because charging facilities have physical limits but also that certain requests are simply not economically attractive. How to make admission decisions is not trivial. The idea of admission control in the context of deadline scheduling takes the form of “customer notification” [21] where the service provider can choose to accept or decline a customer’s request upon the request arrival; see also [22, 23] and references therein. The optimal admission control is not known in the literature.

II. OPTIMAL DEADLINE SCHEDULING

A. Problem Formulation

A job $J = (r, l, d, v)$ is represented by a quadruple specified by the release time r , the processing time l , its deadline d , and the value v . We adopt the so-called proportional value model [24] where the value v of a job is proportional to its processing time l . Without loss of generality, we set $v = cl$, where c is the price per unit charging time. Preemption is allowed at no cost, *i.e.*, a preempted job can be resumed from the point of preemption at a later time.

In the deterministic setting of job scheduling, an input instance I of size n for the scheduler includes n jobs, $I = (J_1, \dots, J_n)$. The input instances in general have different sizes. The optimization is over the collection \mathcal{I} of all finite input instances.

An online scheduler $\mathcal{S}_{\text{online}}$ knows the parameters of job J_i only at its release time r_i . The deadline of associated with each job is firm, *i.e.*, the value of a job can only be collected if it is completed before its deadline. The admission decision is made

immediately, *i.e.*, upon the release of each job, the scheduler has to decide whether to accept or decline the request.

We assume, without loss of generality, that the charging station has access to zero variable cost renewable energy. By relaxing the constraint on the physical capacity on the number of parking spots available, the bottleneck to optimizing profit arises from the limited renewable energy to support all requests; operating cost from purchasing electricity is unavoidable, since all requests of EVs admitted to the facility have to be met.

The operating profit is the difference of the revenue generated from payments of EV customers and the cost of purchasing electricity from the grid. Given an instance I , we denote by $S_{\text{online}}(I)$ the total profit obtained by the online scheduler $\mathcal{S}_{\text{online}}$ and $S_{\text{offline}}(I)$ the profit of the offline scheduler $\mathcal{S}_{\text{offline}}$. In contrast to the online scheduler, an offline scheduler $\mathcal{S}_{\text{offline}}$ knows the entire input instance a priori. We denote by $\mathcal{S}_{\text{offline}}^*$ the optimal offline scheduler.

Our objective is to make the online scheduler competitive across all instances in \mathcal{I} . The measure of competitiveness is defined as follows.

Definition 1. Competitive ratio: An online algorithm $\mathcal{S}_{\text{online}}$ is α -competitive for an input set \mathcal{I} if

$$\min_{I \in \mathcal{I}} \frac{S_{\text{online}}(I)}{S_{\text{offline}}^*(I)} \geq \alpha \quad (1)$$

where I varies over all possible input instances in \mathcal{I} .

In other words, an α -competitive online algorithm is guaranteed to achieve at least α fraction of the optimal offline value under any input instance I in the input set \mathcal{I} .

For the rest of the paper, the input set \mathcal{I} is fixed to be the set of all input instances I with finite sizes, and each service request is valid, *i.e.*, the request can be met if the EV is charged immediately ($d_i \geq r_i + l_i$).

B. Optimal Competitive Ratio

We state the result of the optimal competitive ratio in Theorem 1 followed by a detailed description of the online scheduling algorithm (TAGS) that achieves the optimal competitive ratio.

Theorem 1. Let M be the number of chargers powered by the zero-cost renewable energy. Let c be the price of providing charging service measured by dollars per kilo-Watt-hour (\$/kWh), p (\$/kWh) the marginal cost of electricity from the grid, and $\gamma = c/p$ be the normalized charging price. For the collection \mathcal{I} of all finite job instances, the maximum achievable competitive ratio is given by

$$C^* = \begin{cases} 1 & \text{if } \gamma \geq 1; \\ \frac{1}{z(M, \gamma)} & \text{if } \gamma < 1. \end{cases} \quad (2)$$

where $z(M, \gamma) > 1$ is the unique positive real root larger than 1 of the polynomial $f_{M, \gamma}(z) = \gamma(z + M)^{M+1} - z(M + 1)^{M+1}$.

Proof Sketch: The proof of this theorem involves showing (2) is an upper bound on competitive ratio. To this end, we

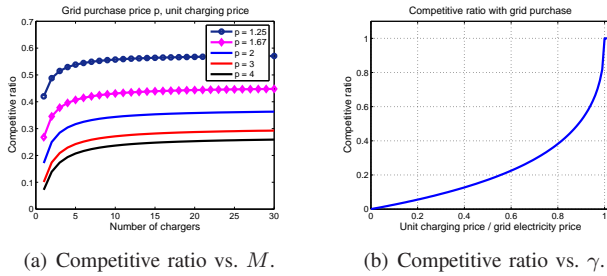


Fig. 2. Competitive ratio vs the number of renewable chargers and normalized charging price.

use an adversary argument by considering a game between the online scheduler and its offline counterpart. The upper bound is then shown to be tight by considering a specific scheduling algorithm TAGS, which is shown in the next section. See [1] for the details of the proof. \square

Theorem 1 divides the normalized price γ into two operating regimes. For the high price regime, the optimal online scheduler is as good as the optimal offline scheduler. Intuitively, when $\gamma > 1$, the price of charging is higher than the price of electricity from the grid; profit is always guaranteed in every job. Thus every request should be accepted, *i.e.*, there should be no admission control.

For the low price regime, $\gamma < 1$, and the maximum achievable competitive ratio is less than one. While the optimal competitive ratio is not intuitive, it is expected that C^* should be monotone with respect to the normalized price γ . This is confirmed in Figure 2(b) (for $M = 1$) where we see a sharper increase as γ is approaching to 1.

When the number of chargers powered by the renewable increases, most EVs can be charged by renewable energy. Thus higher the M is, the less energy is needed from the grid, the less valuable the knowledge about future arrivals is, and the less advantage the optimal offline algorithm has over the best online algorithm. We then expect that the optimal competitive ratio increases with M . This too is confirmed by Figure 2(a).

It is interesting to note from Figure 2(a) that, for any fixed price $\gamma < 1$, the optimal competitive ratio does not approach 1 as $M \rightarrow \infty$. This is because the collection of input instances always requires much more energy than the available renewable supply.

C. TAGS: Optimal Scheduling Algorithm

We now focus on developing an algorithm (TAGS) that achieves the optimal competitive ratio given in Theorem 1. The skeleton of TAGS is given in Figure 3. Here we provide an abbreviated narrative for its implementation as there are many variations that can be incorporated.

The scheduling aspect of TAGS is in fact similar to EDF or LLF. The classical EDF and LLF do not use admission control. It is obvious that it will perform badly when the arrival rate of EVs is heavy. It is however not difficult to modify EDF and LLF using a naive admission control where, for each arrival, the scheduler checks whether the system remains underloaded. Such a check only has low complexity. The

resulting EDF or LLF, referred to as conservative EDF/LLF, completes all charging using renewable only. We can view TAGS as an optimal tradeoff between the conservative and aggressive versions of EDF/LLF.

If a newly released request can be accommodated by a certain charger without affecting any previously accepted requests on this charger, TAGS algorithm will accept the request and add the request to the job queue of the specific charger.

Otherwise, there is no charger that can accommodate the request without affecting any previously accepted requests. The basic structure of TAGS can be broken down into a two-stage decision process for such a “difficult” request. The first is an admission decision that determines whether the arrival request should be accepted to the host charger; the host charger is the charger currently intended for the requests that requires a profitability test. The index of the host charger will initially be set to be 1 and every time a “difficult” request is accepted to the host charger, the index of the host charger increments by 1.

To this end, TAGS performs a threshold test on profitability for the “difficult” request. Specifically, if the accommodation of a request increases the cost of some existing charging orders, TAGS compares the profit of accepting the request and that of declining it. In particular, if the potential profit of acceptance is $(1 + \beta)^{1/M}$ greater than that of declination, the EV request is accepted. Otherwise, the request is declined. Here β is a parameter to be optimized. It should be noted that the profit ratio is a function of the number M of chargers powered by the renewable. It is obvious that, as M increases, TAGS accepts more EV charging orders.

Once a “difficult” request is accepted, TAGS assigns the charging request to the current host charger. Once the request is assigned to a charger, a greedy scheduling of earliest deadline first is used to determine the request to be executed. The scheduler of TAGS maintains a job queue and a tentative earliest deadline first schedule for each processor at all times.

Theorem 2. TAGS achieves the maximum competitive ratio C^* in (2) in Theorem 1.

Proof Sketch: The proof of the optimality of TAGS involves several steps. From Theorem 1, we already have an upper bound for competitive ratio. Here we analyze the performance of TAGS by partitioning jobs into several types and analyze the performance. Specifically, we partition the entire request sets into the *constrained-ended* and *free-ended* sets. The optimal offline value from the *constrained-ended* set is upper bounded by the length of the busy interval for the TAGS algorithm, while the optimal value from the *free-ended* set is upper bounded by the *free-ended* value collected by TAGS algorithm. See [1] for details. \square

III. SIMULATION RESULTS

We present in this section numerical simulation results aimed at providing an assessment of the *average performance* of TAGS. It is not unusual that a competitive ratio optimal algorithm performs poorly when it is measured by an average

TAGS procedure

```
1: set host processor  $k = 1$ 
2: loop
3:   upon event: job  $J_{\text{arr}}$  is released
4:   if  $J_{\text{arr}}$  can be accommodated at some processor  $1 \leq j \leq M$ 
     then
5:     add  $J_{\text{arr}}$  to the job queue of processor  $j$ 
6:   else
7:     examine total Profit-of-accept and Profit-of-decline for
       host processor  $k$ 
8:     if Profit-of-accept  $> (1 + \beta)^{1/M} \times$  Profit-of-decline then
9:       add  $J_{\text{arr}}$  to the job queue of processor  $k$ 
10:      reduce the processing time of  $J_{\text{arr}}$  accordingly
11:      increment host processor  $k$ 
12:     else
13:       decline  $J_{\text{arr}}$ 
14:     end if
15:   end if
16: end loop
```

Fig. 3. TAGS algorithm.

performance metric. Here we evaluate the average profit of TAGS compared with standard benchmarks.

A. Benchmarks and Performance Measure

The first benchmark is the unmanaged charging (UC) which mimics the operations in self-service gas stations. Specifically, an EV is charged immediately until the request is fulfilled. The charger uses renewable if available. Otherwise, it purchases the electricity. No admission policy is used.

The second benchmark is a variation of EDF. Since the classical EDF performs poorly without admission control, we consider the conservative EDF where we equip the classical EDF with a conservative admission control policy that admits a job only if it can be finished. A version of multiprocessor EDF is used in the simulation.

The average profit of the CSP is the metric of evaluation. With energy of use (EOC) pricing, the CSP receives payment from individual EV customer at the price of c \$/kWh. Besides the charging price, the EV traffic also affects the overall profit of CSP, both in term of of traffic intensity and traffic statistics.

When evaluating the profit against traffic intensity, there are two different scenarios. One is treating traffic as an *exogenous* input. In this case, we vary the traffic intensity and evaluate how profit varies. In this case, the traffic includes only those who have decided to charge regardless the charging price. The more significant evaluation is to use *endogenous* traffic in which the traffic includes all EV customers whose decision of charging at a facility depends on the posted price c . In this paper, we present the results correspond to the latter.

B. Pricing Considerations

The endogenous traffic is the result of pricing induced customer trimming for the charging operations. The customer trimming summarizes the effect that the pricing mechanism, together with the customer response, shapes the fraction of

customers who accept the offered price. It is this fraction that determines the endogenous traffic intensity.

The impact on the portfolio of charging energy consumed is directly related to the marginal cost of the charging facility, since more renewable energy implies lower cost, and vice versa. Specifically, if the renewable energy availability is fixed, when the overall charging load is small, with proper scheduling the majority of the charging can be fulfilled by the cheap renewable energy and the marginal cost is low. When the overall charging load increases, the component of the relatively expensive grid electricity inevitably expands and the marginal cost will increase as a result.

Typically, the charging price has to be above the marginal cost, although there is no explicit formula for the marginal cost available for the charging facility with deadlines. Along this line, the pricing mechanism has to strike reasonable balance: a high charging price turns away too many customers whereas a low price pushes up the charging load and causes the increase of marginal cost due to increased purchase of electricity.

C. Simulation Setup

The traffic parameters in the simulation are adopted as follows: the customer arrival process is assumed to be Poisson process with mean inter-arrival time $\lambda^{-1} \in \{0.5, 1, 2\}$ minutes, the charging time requirements assumed to be i.i.d. uniform in the interval $[0, 30]$ minutes, and the relative deadlines assumed to be i.i.d exponential with mean 40 minutes, where the relative deadline is $d - r - l$, the difference between the time span of customer parking $d - r$ and the requested charging time l .

The pricing function simulated is EOU pricing with unit price $c \in [0.03 : 0.01 : 0.27]$ \$/kWh. We conduct Monte Carlo runs with time duration of 8 hours and 6 renewable chargers, each of which is capable of fully charging an EV in 30 minutes. The grid electricity price is set to be constant during the eight hour time frame with unit grid purchase price $p=0.16$ \$/kWh (we raise the grid electricity price from the data entry 0.1132 \$/kWh in EIA Monthly Energy Review for July 2011¹, since there is concern on the high peak powered needed by the charging facilities, thus an incentive to raise the grid purchase price for charging facilities to mitigate the impact of the spiky peak power needed). The customer response curve $\theta(u)$ is assumed to be a sigmoid function $\theta(u) = \frac{\exp(-45(u-0.17))}{(1+\exp(-45(u-0.17)))}$, the shape of which is shown in Fig. 4.

D. Profit Comparisons and Optimal Pricing

The monopolistic profit is plotted in Fig. 4 versus the normalized unit charging price ($\gamma = c/p$) under EOU pricing mechanism. Several observations can be made from Fig. 4: 1) there exists a unique unit price which corresponds to the maximum monopolistic profit, which suggests a relation between the EV adoption level and optimal pricing. 2) TAGS

¹Table 9.9, Transportation sector, Monthly Energy Review, United States Energy Information Administration (EIA), May 2012.

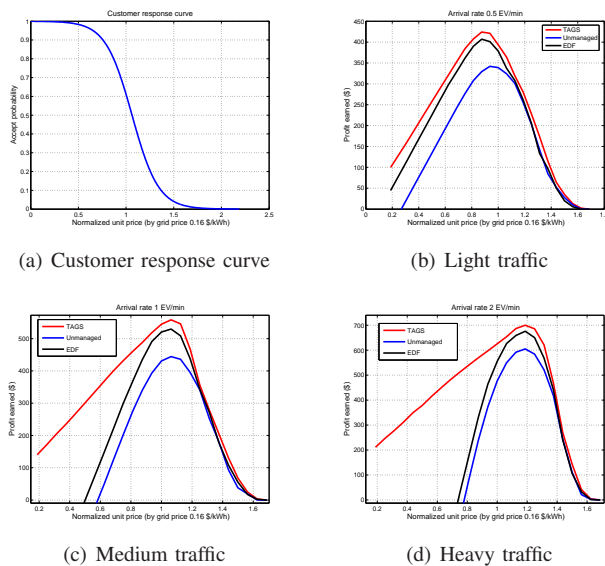


Fig. 4. Impact of scheduling algorithms on operator profitability

outperforms EDF and UC in terms of the maximum monopolistic profit 3) the best unit price increases with the arrival rate (with a smaller arrival rate there are less customers, and the unit price has to be reduced to attract more business so that the charging station capacity can be maintained at a high level) 4) for the same unit price the profit per vehicle decreases when the arrival rate increases, which demonstrates the increasing trend of the marginal cost with the overall charging workload.

IV. CONCLUSION

We have considered in this paper an iEMS for the charging of EVs in large parking facilities using a combination of renewable and purchased energy. The proposed architecture for iEMS is based on the concept of network switched charging, which provides a hardware-software platform that optimally manages a large number of EV customers.

The charging scheduling algorithm (the central controller of iEMS) optimizes jointly the procurement of energy, the charging of EVs, and the admissions of new arrivals. It should be noted that the optimal online scheduling algorithm (TAGS) has very low computation cost. In the simulations, despite that TAGS is optimized for the worst case performance, TAGS performs better than benchmarks when measured by statistical averages.

Several generalizations are currently being considered. Note that the general principle of TAGS, threshold admission (or assignment) and greedy scheduling, can be applied to many new scenarios, including the integration of local storage and variable pricing of purchased electricity. Mechanisms for pricing are also being studied [25].

REFERENCES

[1] S. Chen and L. Tong, "Online deadline scheduling with customer notification: Optimal competitive ratio," in preparation.

[2] M. Caramanis and J. M. Foster, "Management of electric vehicle charging to mitigate renewable generation intermittency and distribution network congestion," in *Proceedings of the 48th IEEE Conference on Decision and Control*, Dec 2009, pp. 4717–4722.

[3] M. Kefayati and C. Caramanis, "Efficient energy delivery management for PHEVs," in *2010 First IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Oct 2010, pp. 525–530.

[4] C. Pang, P. Dutta, S. Kim, M. Kezunovic, and I. Damnjanovic, "PHEVs as dynamically configurable dispersed energy storage for V2B uses in the smart grid," in *7th MedPower*, Nov 2010, pp. 1–6.

[5] L. Gan, U. Topcu, and S. H. Low, "Optimal decentralized protocol for electric vehicle charging," in *Proceedings of Conference on Decision and Control*, 2011.

[6] S. Sojoudi and S. H. Low, "Optimal charging of plug-in hybrid electric vehicles in smart grids," in *Proc. IEEE PES General Meeting*, 2011.

[7] Z. Ma, D. S. Callaway, and I. A. Hiskens, "Decentralized charging control of large populations of plug-in electric vehicles," *IEEE Trans. Control Systems Technology*, 2011.

[8] P. Kulshrestha, L. Wang, M.-Y. Chow, and S. Lukic, "Intelligent energy management system simulator for phev at municipal parking deck in a smart grid environment," in *IEEE PES General Meeting 2009*, 2009.

[9] W. Su and M.-Y. Chow, "Performance evaluation of a phev parking station using particle swarm optimization," in *IEEE Power and Energy Society General Meeting 2011*, Detroit, MI, USA, July 2011.

[10] A. Subramanian, M. Garcia, A. Dominguez-Garcia, D. Callaway, K. Poolla, and P. Varaiya, "Real-time scheduling deferrable electric loads," in *2012 American Control Conference*, 2012.

[11] S. Chen, T. He, and L. Tong, "Optimal deadline scheduling with commitment," in *2011 Allerton Conference on Communication, Control and Computing*, Oct 2011.

[12] S. Chen, Y. Ji, and L. Tong, "Deadline scheduling for large scale charging of electric vehicles with renewable energy," in *2012 IEEE Sensor Array and Multichannel Signal Processing Workshop - Special Session on Smart Grid*, June 2012.

[13] —, "Large scale charging of electric vehicles," in *Proceedings of IEEE PES General Meeting*, 2012.

[14] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *Journal of ACM*, vol. 20, pp. 46–61, 1973.

[15] M. Dertouzos, "Control robotics: the procedural control of physical processes," in *International Federation for Information Processing Congress 1974*, Stockholm, Sweden, August 1974, pp. 807–813.

[16] A. Mok, "Fundamental design problems of distributed systems for the hard real-time environment," Ph.D. dissertation, MIT, 1983.

[17] C. D. Locke, "Best-effort decision-making for real-time scheduling," Ph.D. dissertation, CMU, 1986.

[18] S. Baruah, G. Koren, D. Mao, B. Mishra, A. Raghunathan, L. Rosier, D. Shasha, and F. Wang, "On the competitiveness of on-line real-time task scheduling," *Real-Time Systems*, vol. 4, pp. 125–144, 1992.

[19] G. Koren and D. Shasha, "Dover: An optimal on-line scheduling algorithm for overloaded uniprocessor real-time systems," *SIAM Journal of Computing*, vol. 24, pp. 318–339, 1995.

[20] —, "MOCA: A multiprocessor on-line competitive algorithm for real-time system scheduling," *Theor. Comput. Sci.*, pp. 75–97, 1994.

[21] A. Bar-Noy, J. A. Garay, and A. Herzberg, "Sharing video on demand," *Discrete Applied Mathematics*, vol. 129, no. 1, pp. 3–30, 2003.

[22] N. Thibault and C. Laforest, "Online time constrained scheduling with penalties," in *23rd IEEE International Symposium on Parallel and Distributed Processing*, 2009.

[23] S. P. Y. Fung, "Online preemptive scheduling with immediate decision or notification and penalties," in *COCOON'2010*, 2010.

[24] M. Goldwasser and B. Kerbikov, "Admission control with immediate notification," *Journal of Scheduling*, vol. 6, pp. 269–285, 2003.

[25] S. Chen, T. Mount, and L. Tong, "Pricing and competition for large scale charging of electric vehicles," in *Hawaii Intl. Conf. on System Sciences*, 2013, submitted.